

On Small Area Prediction Interval Problems

Snigdhasu Chatterjee, Parthasarathi Lahiri, Huilin Li
University of Minnesota, University of Maryland, University of Maryland

Abstract

Empirical best linear unbiased prediction (EBLUP) method uses a linear mixed model in combining information from different sources of information. This method is particularly useful in small area problems. The variability of an EBLUP is traditionally measured by the mean squared prediction error (MSPE), and interval estimates are generally constructed using estimates of the MSPE. Such methods have shortcomings like under-coverage or over-coverage, excessive length and lack of interpretability. We propose a parametric bootstrap approach to estimate the entire distribution of a suitably centered and scaled EBLUP. The bootstrap histogram is highly accurate, and differs from the true EBLUP distribution by only $O(d^3 n^{-3/2})$, where d is the number of parameters and n the number of observations. This result is used to obtain highly accurate prediction intervals. Simulation results demonstrate the superiority of this method over existing techniques of constructing prediction intervals in linear mixed models.

Keywords: Predictive distribution, prediction interval, linear mixed model, small area, bootstrap, coverage accuracy.

1 Introduction

Large scale sample surveys are usually designed to produce reliable estimates of various characteristics of interest for large geographic areas. However, for effective planning of health, social and other services, and for apportioning government funds, there is a growing demand to produce similar estimates for smaller geographic areas and for other sub-populations. To meet this demand, it is necessary to supplement the survey data with other relevant information that is often obtained from different administrative and census records. In many small area applications, linear mixed models are now routinely used in combining information from various sources and explaining different sources of errors. These models incorporate area specific random effects which explain the “between small area variations”, not otherwise explained by the fixed effects part of the model.

For a good review on small area and linear mixed model research, the readers are referred to the book by Rao (2003), and two recent review papers by Rao (2005) and Jiang and Lahiri (2006). Several other applications of linear mixed models may be found in McCulloch and Searle (2001). Point prediction using the empirical best linear unbiased predictor (EBLUP) and the associated mean

squared prediction error (MSPE) estimation have been studied extensively. See Jiang, Lahiri and Wan (2002), Rao (2005) and Jiang and Lahiri (2006) for a review on the subject, especially on the latest development on resampling methods for MSPE estimation. However, little progress has been made outside the basic study of the first two moments, for example on the properties of quantiles (central or tail) of predictors, or on the effect of high dimensionality of the parameters.

For example, research on interval estimates in small area studies is typically limited to special cases of the Fay-Herriot model (described in detail in Section 2), where the traditional estimates are of the form $EBLUP \pm z_{\alpha/2} \sqrt{mspe}$. Here $mspe$ is an estimate of the true $MSPE$ of the EBLUP, and $z_{\alpha/2}$ is the upper $100(1-\alpha/2)\%$ point of the standard normal distribution. The coverage probabilities of such intervals may converge to the nominal level $1-\alpha$; but the intervals are not efficient since they have either under-coverage or over-coverage problem, depending on the particular choice of the MSPE estimator. More precisely, the coverage error of such interval is of order $O(n^{-1})$ or higher, which is not accurate enough for most applications of small area studies with small sample size n .

In this paper we address the problem of approximating the distribution of a predictor, and applying it to obtain prediction intervals, in a very general framework of linear mixed models. We consider the following model from Das, Jiang and Rao (2004):

$$\mathbf{Y}_n = \mathbf{X}\beta + \mathbf{Z}\mathbf{v}_q + \mathbf{e}_n, \quad (1.1)$$

where $\mathbf{Y}_n \in \mathbb{R}^n$ is a vector of observed responses; $\mathbf{X}_{n \times p}$ and $\mathbf{Z}_{n \times q}$ are known matrices; \mathbf{v}_q and \mathbf{e}_n are independent random variables with dispersion matrices $D_q(\psi)$ and $R_n(\psi)$ respectively. Here $\beta \in \mathbb{R}^p$ and $\psi \in \mathbb{R}^k$ are fixed parameters.

The mixed ANOVA and the longitudinal models, including the Fay-Herriot and the nested error regression models, are special cases of (1.1). We can consider both balanced and unbalanced lay-outs in the above framework. In addition, we develop our theory and methodology that allow the parameter dimension $d = p + k$ to grow with sample size n . Dimension dependent asymptotics are extremely important in the current context, since many small areas may have sample sizes comparable to dimensions of the regression and variance components parameters; see, for example, Jiang (1996) for their use and importance in linear mixed models.

We employ a parametric bootstrap method to approximate the distribution of a predictor. We concentrate on

the *empirical best linear predictor* (EBLUP), owing to its wide popularity and use. In Theorem 3.1 we show that the bootstrap histogram incurs an error of $O(d^3 n^{-3/2})$ in approximating the distribution of a centered and scaled EBLUP. For estimating the distribution of centered and scaled *estimators*, under standard regularity conditions and fixed d , the normal approximation based on the Central Limit Theorem has an error of $O(n^{-1/2})$, but the bootstrap can achieve higher order accuracy, with typical approximation error of $O(n^{-1})$. Theorem 3.1 may be seen as an extension of this higher order accuracy phenomenon, in the context of prediction. Although our motivation and terminology comes from small area context, our bootstrap methodology and theoretical results are directly applicable to other usages of linear mixed models.

There are several potential applications of a highly accurate approximation of the entire distribution of the EBLUP. For example, it may be used to obtain (a) bagging predictors, (b) computing mean squared errors or other risks, (c) hypothesis testing, (d) calibration of traditional estimators, and (e) prediction interval construction. In this paper, we concentrate on the last application, since prediction intervals combine features of both point prediction and hypothesis testing nicely, and have not been extensively explored in small area or other linear mixed model contexts.

Prediction intervals are useful in small area studies in several ways. For example, prediction intervals may help establish if different counties have similar resources and needs, or if different ethnic or other sub-population groups are equally exposed to a particular disease. Our simultaneous concentration on dimension asymptotic is also relevant. It has long been recognized that health, economic activity and other measures of human well-being depend on a number of exogenous and endogenous factors, many of which must be measured at the individual level and incorporated in the model. In statistical terms, this translates to high dimensionality of β and ψ .

In Section 2, we review a few existing techniques for predictor distribution approximation and interval estimate construction. We pay special attention to the usage of resampling in such approximations/constructions. For prediction intervals, available literature is heavily concentrated on special cases of the Fay-Herriot model. Since traditional intervals perform poorly in terms of coverage or length or both, many attempts have been made to fine tune and calibrate them, often using resampling. To the best of our knowledge, approximation of the entire distribution of a predictor has not been attempted in general small area problems, and we briefly review the related research for independent data.

In Section 3, we present our bootstrap algorithm for model (1.1). Our main result states that the sup-norm distance between the distribution of EBLUP and its bootstrap approximation is $O(d^3 n^{-3/2})$. A direct corollary is that the bootstrap prediction interval has coverage accuracy of $O(d^3 n^{-3/2})$. Note that our proposed prediction interval is a *bootstrap interval*, which is different from the

traditional approaches of obtaining asymptotic intervals first and then calibrating it. Our interval can be calibrated one or more times to achieve a coverage accuracy of $O(d^5 n^{-5/2})$ or higher, if needed.

We performed several simulation experiments in order to study how our percentile bootstrap interval estimate compares with existing techniques. A sample of these studies are reported in Section 4. The main message from the simulations is that the prediction intervals resulting from the proposed parametric bootstrap perform considerably better than the traditional techniques, which is a reflection of the high order accuracy theoretically established in Section 3. The proofs of all the theorems will appear in Chatterjee, Lahiri and Li (2007).

2 A review of predictor distribution approximation and interval construction

2.1 Approximating distributions of predictors

Considerable theoretical research has been carried out in the prediction of a random variable that is independent of \mathbf{Y}_n , and has density $\xi(\cdot|\beta, \psi)$. In terms of expected Kullback-Leibler divergence, the naive plug-in predictor density $\xi(\cdot|\hat{\beta}, \hat{\psi})$ performs poorly compared to Bayesian predictors $\int \xi(\cdot|\beta, \psi) d\pi(\beta, \psi|\mathbf{Y}_n)$, where $\pi(\beta, \psi|\mathbf{Y}_n)$ denotes the posterior density of (β, ψ) with respect to a suitable prior; see, for example, Aitchison (1975), Murray (1977), Ng (1980), Komaki (1996, 2001, 2006), and George *et al* (2006). Harris (1989) showed that the *bootstrap predictor*

$$\xi^*(\cdot) = \int \xi(\cdot|s, t) d\mathcal{L}^*(s, t), \quad (2.2)$$

where \mathcal{L}^* denotes the bootstrap approximation to the distribution of $(\hat{\beta}, \hat{\psi})$, also performs better than the naive plug-in predictor. Fushiki, Komaki and Aihara (2004) showed that the bootstrap predictor (2.2) is asymptotically equivalent to a Bayesian predictor with Hartigan's M -prior. The M -prior has certain optimality properties which may be found in Hartigan (1964, 1998). In a related work, Fushiki, Komaki, Aihara (2005) showed that the Harris predictor is related to bagging of Breiman (1996).

In small area or other linear mixed model contexts, the random variable of interest depends on \mathbf{Y}_n , unlike the framework described above. Also, performance measures other than expected Kullback-Leibler divergence may be of interest, for example, the length and coverage of prediction intervals.

2.2 A review of interval estimation techniques

For a general linear mixed model, Jeske and Harville (1988) proposed a prediction interval for a mixed effect, but did not include the effect of estimated unknown variance components on the accuracy of their proposed interval.

Jiang and Zhang (2002) used a distribution-free method for constructing prediction intervals for a future observation under a non-Gaussian linear mixed model, based on the theory developed by Jiang (1998). This technique does not employ any area specific information and can be useful in constructing intervals when there is no survey data on the response variable. Jiang and Zhang (2002) proposed another method which can be applied to the situation when the sample size is large within each area. This is a technique of first obtaining the EBLUP for the random effects and the residuals. Then, under conditions sufficient to imply that the number of times each random effect is repeated (*i.e.*, number of observations in each small area) tends to infinity, the empirical distribution of random effects as well as the residuals converge appropriately. This technique fails when we do not have large samples for each small area, a situation that is common in many small area applications.

Recently, Hall and Maiti (2006b) have studied a parametric bootstrap for mixed models in several aspects, including interval estimation. A review of their approach towards interval estimation may be found in Rao (2005). In Section 3, we discuss in detail how their model, results and asymptotics differ from ours.

Other than the above three papers, research on small area prediction intervals is largely concentrated on special cases of the Fay-Herriot model, described below:

1. Conditional on $\theta = (\theta_1, \dots, \theta_n)^T$, $\mathbf{Y}_n = (Y_1, \dots, Y_n)^T$ follows a n -variate normal distribution with mean θ and dispersion matrix D with *known* diagonal entries $D_i > 0$ and off-diagonal entries 0. Here (and in the sequel) all vectors are taken to be column vectors; for any vector (matrix) a (A), the notation a^T (A^T) denotes its transpose.
2. The variable θ follows a n -variate normal distribution with mean $\mathbf{X}\beta$ for a known $n \times p$ matrix \mathbf{X} and unknown but fixed vector $\beta \in \mathbb{R}^p$. The dispersion matrix is $A\mathbf{I}_n$, where the matrix \mathbf{I}_n is the n dimensional identity matrix and A is an unknown constant.

There are several options for constructing interval estimates for $\theta_i = \mathbf{x}_i^T \beta + v_i$. One may use only the Level 1 model for the observed data, or only the Level 2 model for the borrowed strength component, or a combination of both. The interval for θ_i based only on the Level 1 model is given by $I_i^D(\alpha) : Y_i \pm z_{\alpha/2} D_i^{1/2}$. Obviously, for this interval, the coverage probability is $1 - \alpha$. However, it is not efficient, since its average length is too large to make any reasonable conclusion. This is due to the high variability of the point predictor Y_i .

An interval based only on Level 2 ignores the crucial area specific data that is modeled in Level 1, and hence falls short on two counts: it fails to be relevant to the specific small area under consideration, and it fails to achieve sufficient coverage accuracy. A simple example given later in this section demonstrates this latter property.

Thus, interval estimation techniques that combine both levels of the Fay-Herriot model are required. A popular approach is to employ empirical Bayes methodology. Cox (1975) proposed the following empirical Bayes interval:

$$I_i^C(\alpha) : (1 - \hat{B}_i)Y_i + \hat{B}_i \mathbf{x}_i^T \hat{\beta} \pm z_{\alpha/2} D_i^{1/2} (1 - \hat{B}_i)^{1/2}$$

where \hat{B}_i and $\hat{\beta}$ are estimators of $B_i = D_i/(A+D_i)$ and β respectively, and \mathbf{x}_i is the i^{th} row of \mathbf{X} . Under standard regularity conditions, $\mathbb{P}(\theta_i \in I_i^C(\alpha)) = 1 - \alpha + O(n^{-1})$, where \mathbb{P} denotes a probability measure induced by the joint distribution of Level 1 and Level 2. Thus, this prediction interval attains the desired coverage probability asymptotically, but the coverage error is of order $O(n^{-1})$, which is not accurate enough for many small area applications. This lack of accuracy may partially be due to the additional variability resulting from estimation of β and A . Currently, second-order or nearly unbiased MSPE estimators are available in several linear mixed models, see for example, Jiang *et al* (2002), Datta, Rao and Smith (2005), Hall and Maiti (2006a). Naive empirical Bayes intervals constructed using EBLUP, MSPE estimators and standard normal quantiles typically have an error of $O(n^{-1})$ or higher.

For a special case of the Fay-Herriot model with common mean and equal sampling variances $D_i = D$, Morris (1983a) incorporated the additional uncertainty due to the estimation of the hyperparameters. However, Basu, Ghosh and Mukerjee (2003) showed that the resulting empirical Bayes interval proposed by Morris (1983a) still has coverage error of $O(n^{-1})$. They used analytical calibration of the Morris' interval to reduce the coverage error to $o(n^{-1})$. They also showed that with suitable analytical approximations in place, an interval due to Carlin and Louis (1996, p.98), and a new interval, have coverage error of the order $o(n^{-1})$. Datta, Ghosh, Smith and Lahiri (2002) used similar analytical calibration in a more general Fay-Herriot model, and obtained a prediction interval with coverage error of $O(n^{-3/2})$. Morris (1983b) considered a variation of his (1983a) work with the use of a hierarchical Bayes type point estimator. Hill (1990) suggested a general framework, which in the Fay-Herriot setting, matches with an exact hierarchical Bayes confidence interval. Datta *et al.* (2002) followed up Hill's idea to obtain an interval with coverage error of $O(n^{-1})$.

Apart from the analytical approaches, calibration using different bootstrap techniques has been popular. The methods differ in the generation of the bootstrap samples and the type of correction made. For a special case of the Fay-Herriot model where Y_1, \dots, Y_n are independent and identically distributed, Laird and Louis (1987) proposed three different resampling strategies: (a) usual nonparametric bootstrap with replacement with the data, (b) a semi-parametric method, assuming density at the first level of their two level model is known but that at the second level is unknown, and (c) a parametric bootstrap. In linear mixed models, the nonparametric and semi-parametric bootstrap approximation of the distri-

bution of the EBLUP are generally not consistent. Once the bootstrap sample (nonparametric, semi-parametric or parametric) is generated, the next challenge is to find a method that corrects the empirical Bayes confidence intervals $I_i^C(\alpha)$ to achieve better coverage. Laird and Louis (1987) considered an imitation of the hierarchical Bayes approach.

Carlin and Gelfand (1990, 1991) point out that the hierarchical Bayesian methods like those of Laird and Louis (1987) lead to a *lengthening* of the empirical Bayes interval, which is not the same as a *correction*. They discuss an example where increasing the length further exacerbates the coverage bias. They suggest parametric bootstrap to calibrate the empirical Bayes interval.

Calibration of intervals has been one of the major uses of bootstrap for some time, and can lead to considerable improvement of coverage accuracy. Coupled with use of bias correction, use of pivotal or nearly pivotal statistics, and Edgeworth corrections, improvements from calibration can sometimes be dramatic. See Abramovitch and Singh (1985), Beran (1990a, 1990b), the book by Efron and Tibshirani (1993) and references therein for further details on these issues. On the other hand, calibration is both time and computational effort intensive, often requiring iterative searches; it typically increases variability; the results often lack straightforward interpretability; and successive calibrating steps typically have diminishing returns in terms of improvement of coverage. It is not always clear what property of an interval, *ie*, length, coverage, end points or some other characteristic, ought to be calibrated, see for example DiCiccio and Efron (1996) and the discussions of it by Hall and Martin (1996), Lee and Young (1996); and the interesting example in the rejoinder. Some calibrating options do not exist for multivariate confidence or prediction regions. Asymptotic results suggest calibrated intervals have better coverage accuracy, but do not consider the variability induced by the calibration, do not represent performance in finite samples; or reflect the degree in which the finite sample results depend on unknown parameters and their estimators. Nevertheless, calibration is an excellent tool to improve coverage of intervals; though it seems sensible to use a more accurate interval and little or no calibration; rather than a less accurate interval with intensive calibration. The bootstrap interval we obtain in Section 3 is one such highly accurate interval, and requires the same amount of computational effort as one round of bootstrap based calibration of Carlin and Gelfand (1990, 1991).

Hall (2006) suggested an application of the nonparametric bootstrap confidence interval based on the generated θ_i^* 's only. In the small area context, this may be applicable when the differences between the small areas are minor, or carried only in the fixed effects. In surveys, robustness is always an important issue, and the practitioners are always interested in efficient nonparametric methods. However, due to scarce data at the small area level, nonparametric estimators tend to under-perform, often severely. This is because the nonparametric models

typically permit the generation of bootstrap histograms based on a synthetic model or the regression model, but do not permit approximation of the conditional distribution of θ_i given the data \mathbf{Y}_n . As a result, the nonparametric bootstrap prediction interval for θ_i is likely to underweight the area specific data. Accurate weighting of the area specific data is important for achieving good coverage properties, as the example below shows. Hall (2006) also pointed out the importance of parametric bootstrap in small area estimation and other related problems.

EXAMPLE. Consider the following special case of the Fay-Herriot model where $D_i \equiv 1$, and $\mathbf{x}_i^T \beta \equiv \mu$. Thus, at Level 1, Y_i 's given the θ_i 's are independently distributed as $N(\theta_i, 1)$ random variables; and at Level 2 the θ_i 's are independent, identically distributed as $N(\mu, \tau^2)$ random variables. The estimators of μ and τ^2 are given respectively by $\hat{\mu} = \bar{Y}$, $\hat{\tau}^2 = \max(0, s^2 - 1)$, where $s^2 = \sum (y_i - \bar{y})^2 / (n - 1)$. Assume $\hat{\tau}^2 > 0$, a condition that is satisfied in many problems. The bootstrap procedure would require us to generate $\theta_i^* \stackrel{iid}{\sim} N[\hat{\mu}, \hat{\tau}^2]$ and $Y_i^* | \theta_i^* \stackrel{iid}{\sim} N[\theta_i^*, 1]$. Then we have $\hat{\mu}^* = \bar{Y}^*$, $\hat{\tau}^{2*} = \max(0, s^{*2} - 1)$ where $s^{*2} = \sum (y_i^* - \bar{y}^*)^2 / (n - 1)$. An obvious Level 2 based bootstrap prediction interval for θ_i that is *not* area specific, is given by

$$\left(\hat{\mu} - t_1 \sqrt{\hat{\tau}^2}, \hat{\mu} + t_2 \sqrt{\hat{\tau}^2} \right), \tag{2.3}$$

where (t_1, t_2) are cutoff points satisfying $\mathbb{P}^*(\hat{\mu}^* - t_1 \sqrt{\hat{\tau}^{2*}} \leq \theta_i^* \leq \hat{\mu}^* + t_2 \sqrt{\hat{\tau}^{2*}}) = 1 - \alpha$.

It can be shown that interval (2.3) has coverage of $1 - \alpha + O(n^{-1/2})$ which makes it consistent, but hardly accurate enough. The lack of accuracy is due to the use of the Level 2 distribution only, so that the Level 1 data Y_i plays no special role in the interval construction.

In Bayesian terminology, the Level 2 of the Fay-Herriot model essentially corresponds to a *prior* on θ_i , while the Level 1 model yields the *likelihood*. Using only the “prior knowledge” (Level 2 distribution) does not even yield consistency in general. However, in some instances using the Level 2 distribution in conjunction with bootstrap can have a calibration effect that obtains $O(n^{-1/2})$ consistency, as shown above.

3 Parametric bootstrap prediction interval for a general linear mixed model

We consider the model:

$$\mathbf{Y}_n = \mathbf{X}\beta + \mathbf{Z}\mathbf{v}_q + \mathbf{e}_n \tag{3.4}$$

where \mathbf{X} is a known $(n \times p)$ matrix; \mathbf{Z} is a known $(n \times q)$ matrix; $\mathbf{Y}_n \in \mathbb{R}^n$ is the vector of observed data; $\beta \in \mathbb{R}^p$ is a fixed but unknown parameter vector; $\mathbf{v}_q \in \mathbb{R}^q$ and $\mathbf{e}_n \in \mathbb{R}^n$ are random variables following the normal distributions $N_q(0, D_q)$ and $N_n(0, R_n)$ respectively. The integer q may depend on n , thus $q \equiv q_n$. Assume the

sequence $\{\mathbf{v}_q\}$ and $\{\mathbf{e}_n\}$ are independent. The first term $\mathbf{X}\beta$ represents the fixed effects, and the second term $\mathbf{Z}\mathbf{v}_q$ the random effects. Thus $\mathbf{X}\beta + \mathbf{Z}\mathbf{v}_q$ constitutes the *signal* component of the observed data, while \mathbf{e}_n is the *noise*. The properties of the signal are of interest, which depend on the unknown parameters β , D_q and R_n . Assume that the $(q \times q)$ matrix D_q and the $(n \times n)$ matrix R_n are known up to a $(k \times 1)$ vector of unknown parameters, thus $D_q = D_q(\psi)$ and $R_n = R_n(\psi)$ for a fixed but unknown $\psi = (\psi_1, \dots, \psi_k)^T \in \mathbb{R}^k$. Note that the dispersion matrix of the observed data \mathbf{Y}_n is given by

$$\Sigma_n = \Sigma_n(\psi) = R_n(\psi) + \mathbf{Z}D_q(\psi)\mathbf{Z}^T.$$

We henceforth drop the n from \mathbf{Y}_n , \mathbf{e}_n , R_n and Σ_n , and q from \mathbf{v}_q and D_q to simplify notations. We take $d = p + k$, the dimension of the parameter space. Let $\theta = (\beta, \psi)$ denote the unknown parameters.

Das, Jiang and Rao (2004) show that several linear mixed models, including Analysis of Variance (ANOVA) models and longitudinal models of both balanced and unbalanced nature are special cases of the the model (3.4). Unbalanced ANOVA models arise, for example, when $R = \sigma_0^2 \mathbf{I}_n$; and $D = \text{diag}(\sigma_1^2 \mathbf{I}_{r_1}, \dots, \sigma_{k-1}^2 \mathbf{I}_{r_{k-1}})$ where \mathbf{I}_r is the $r \times r$ identity matrix. Here ψ is the vector of variance components $\psi = (\sigma_0^2, \dots, \sigma_{k-1}^2)$. Unbalanced longitudinal models arise when Σ has a block diagonal structure.

Let $T = c^T(\mathbf{X}\beta + \mathbf{Z}\mathbf{v})$, where c is a fixed and known $(n \times 1)$ vector. The case where c is a $n \times m$ matrix obtains multidimensional predictive quantities, and their treatment is similar to the univariate case described below, with some minor algebraic variations. We concentrate on univariate T for easier exposition. The conditional distribution of T given \mathbf{Y} is $N(\mu_T, \sigma_T^2)$, where

$$\begin{aligned} \mu_T &= c^T \mathbf{X}\beta + c^T \mathbf{Z}D\mathbf{Z}^T \Sigma^{-1}(\mathbf{Y} - \mathbf{X}\beta) \\ &= c^T R \Sigma^{-1} \mathbf{X}\beta + c^T \mathbf{Z}D\mathbf{Z}^T \Sigma^{-1} \mathbf{Y}, \end{aligned} \quad (3.5)$$

$$\sigma_T^2 = c^T \mathbf{Z} (D - D\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}D) \mathbf{Z}^T c. \quad (3.6)$$

Generally, β and ψ (and hence D and R) are estimated from the data \mathbf{Y} by using the marginal distribution of \mathbf{Y} , given by $N_n(\mathbf{X}\beta, \Sigma)$. The resulting estimates $\hat{\mu}_T$ and $\hat{\sigma}_T$ of the mean and variance of T are expressions similar to (3.5) and (3.6), with $\hat{\beta}$ and $\hat{\psi}$ in place of β and ψ .

For algebraic simplicity, in the rest of this paper we assume that \mathbf{X} is full column rank and use the estimator $\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$. This is the *ordinary least squares* estimator of β . Using other estimators like the maximum likelihood estimator or the *weighted least squares* estimator, with appropriate conditions on the weights, is another possibility. This makes little difference in the asymptotic analysis as long as the weights are smooth functions of ψ . Estimator $\hat{\psi}$ of ψ is typically obtained by maximum likelihood or restricted maximum likelihood techniques.

Based on the fact that $\sigma_T^{-1}(T - \mu_T)$ is a standard Normal pivot, the traditional approach to interval estimation

for T , reviewed in Section 2, is to take $(\hat{\mu}_T \pm z\sqrt{mspe})$ for some estimator $mspe$ of $MSPE$ and the appropriate Normal quantile z . Unfortunately, $\hat{\sigma}_T^{-1}(T - \hat{\mu}_T)$ is not a pivot, and the traditional approach produces too short or too long intervals. Let the distribution of $\hat{\sigma}_T^{-1}(T - \hat{\mu}_T)$ be \mathcal{L}_n . Recognizing that \mathcal{L}_n is not the standard Normal distribution, we propose to estimate it using parametric bootstrap.

Define

$$\mathbf{Y}^* = \mathbf{X}\hat{\beta} + \mathbf{Z}\mathbf{v}^* + \mathbf{e}^*$$

where $\mathbf{v}^* \sim N_q(0, D(\hat{\psi}))$ and $\mathbf{e}^* \sim N_n(0, R(\hat{\psi}))$ are independent of each other.

From \mathbf{Y}^* , obtain $\hat{\beta}^*$ and $\hat{\psi}^*$ using the same techniques used to obtain $\hat{\beta}$ and $\hat{\psi}$ earlier. Next, obtain $\hat{\mu}_T^*$ and $\hat{\sigma}_T^*$ using $\hat{\beta}^*$, $\hat{\psi}^*$, (3.5) and (3.6). Define $T^* = c^T(\mathbf{X}\hat{\beta}^* + \mathbf{Z}\mathbf{v}^*)$. The distribution of

$$\hat{\sigma}_T^{-1*}(T^* - \hat{\mu}_T^*),$$

conditional on the data \mathbf{Y} , is the parametric bootstrap approximation \mathcal{L}_n^* of \mathcal{L}_n . Using this approximation, we then proceed to obtain the interval estimate for T as $(\hat{\mu}_T + q_1 \hat{\sigma}_T, \hat{\mu}_T + q_2 \hat{\sigma}_T)$, where q_1 and q_2 are appropriate quantiles of the bootstrap approximation \mathcal{L}_n^* of \mathcal{L}_n .

Our main result is that \mathcal{L}_n^* approximates \mathcal{L}_n up to $O(d^3 n^{-3/2})$ terms. In order to state the assumptions for our result, let us introduce some terminology and notation now. For any function $f(\psi) : \mathbb{R}^a \rightarrow \mathbb{R}$, $f'(\psi)$ denotes its first derivative written as a $a \times 1$ column vector; $f''(\psi)$ denotes the $a \times a$ second derivative matrix. For a symmetric matrix A , λ_{max} and λ_{min} respectively denote its maximum and minimum eigenvalue.

The following are the assumptions for our result in this section:

1. The following relations hold:

$$\|\mathbf{X}^T c\| = O(1), \quad (3.7)$$

$$\|\mathbf{X}^T \Sigma^{-1} \mathbf{Z}D\mathbf{Z}^T c\| = O(1), \quad (3.8)$$

$$c^T \mathbf{Z}D\mathbf{Z}^T c = O(1), \quad (3.9)$$

$$c^T \mathbf{Z}D\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}D\mathbf{Z}^T c = O(1). \quad (3.10)$$

In addition,

$$\sigma_T^2 = c^T \mathbf{Z} (D - D\mathbf{Z}^T \Sigma^{-1} \mathbf{Z}D) \mathbf{Z}^T c > M > 0,$$

for some constant $M > 0$.

2. Assume that

$$\sup_{1 \leq i \leq n} \sum_{j=1}^p \left[\sum_{a=1}^n X_{ja} \Sigma_{ai}^{1/2} \right]^2 = O(p/n), \quad (3.11)$$

$$\lambda_{min} [n^{-1} \mathbf{X}^T \mathbf{X}] > M > 0, \quad (3.12)$$

for some constant $M > 0$.

3. The eigenvalues of the matrices D and R lie in (L^{-1}, L) for some $L > 1$. The eigenvalues of $D(\hat{\psi})$ and $R(\hat{\psi})$ lie in $(L^{-1}/2, 2L)$. The eigenvalues of Σ lie in a compact set on the positive half of the real line. In the representations

$$D = D_1(\psi)D_1^T(\psi), \quad \hat{D} = D_1(\psi)\Lambda_D(\hat{\psi})D_1^T(\psi), \quad (3.13)$$

$$R = R_1(\psi)R_1^T(\psi), \quad \hat{R} = R_1(\psi)\Lambda_R(\hat{\psi})R_1^T(\psi) \quad (3.14)$$

where Λ_R and Λ_D are diagonal matrices, the following conditions are satisfied: All the entries of the $q \times q$ matrix $\Lambda_D = \text{diag}(\Lambda_{D1}, \dots, \Lambda_{Dq})$ and the $n \times n$ matrix $\Lambda_R = \text{diag}(\Lambda_{R1}, \dots, \Lambda_{Rn})$ have three bounded continuous derivatives. We denote by Λ'_D the $k \times q$ matrix whose $(j, i)^{\text{th}}$ entries are given by

$$((\Lambda'_D))_{j,i}(\psi) = \frac{\partial}{\partial \psi_j} \Lambda_{Di}(\psi).$$

The $(j, i)^{\text{th}}$ entry of the $k^2 \times q$ matrix Λ''_D is

$$((\Lambda''_D))_{j,i}(\psi) = \frac{\partial^2}{\partial \psi_{j_1} \partial \psi_{j_2}} \Lambda_{Di}(\psi).$$

The $(j, i)^{\text{th}}$ entry of the $k^3 \times q$ matrix $\Lambda^{(3)}_D$ is

$$\left((\Lambda^{(3)}_D) \right)_{j,i}(\psi) = \frac{\partial^3}{\partial \psi_{j_1} \partial \psi_{j_2} \partial \psi_{j_3}} \Lambda_{Di}(\psi),$$

where $j_1 + (j_2 - 1)k + (j_3 - 1)k^2 = j$, $j_1, j_2, j_3 = 1, \dots, k$, $j = 1, \dots, k^3$, and $i = 1, \dots, q$. We define the $k \times n$ matrix Λ'_R , the $k^2 \times n$ matrix Λ''_R and the $k^3 \times n$ matrix $\Lambda^{(3)}_R$ along identical lines as above. The following conditions are assumed:

$$\lambda_{max} \Lambda_D^T(\psi) \Lambda'_D(\psi) = O(1), \quad (3.15)$$

$$\lambda_{max} \Lambda_R^T(\psi) \Lambda'_R(\psi) = O(1), \quad (3.16)$$

$$\lambda_{max} \Lambda_D^{''T}(\psi) \Lambda''_D(\psi) = O(1), \quad (3.17)$$

$$\lambda_{max} \Lambda_R^{''T}(\psi) \Lambda''_R(\psi) = O(1), \quad (3.18)$$

$$\lambda_{max} \Lambda_D^{(3)T}(\psi^*) \Lambda^{(3)}_D(\psi^*) < M = O(1), \quad (3.19)$$

$$\lambda_{max} \Lambda_R^{(3)T}(\psi^*) \Lambda^{(3)}_R(\psi^*) < M = O(1), \quad (3.20)$$

for some constant $M > 0$ for all ψ^* in a neighborhood of the true value ψ .

4. Let $S = (k/n)^{1/2}(\hat{\psi} - \psi)$. Assume that all the moments of $\|S\|$ are $O(1)$. Moreover, the following relations are also satisfied:

$$\mathbb{E}S_j = O(\sqrt{k/n}), \quad (3.21)$$

$$\mathbb{E}S_a S_b = O(\sqrt{k/n}), \quad (3.22)$$

$$\mathbb{E}S_j(\mathbf{Zv} + \mathbf{e})_i = O(\sqrt{k/n}), \quad (3.23)$$

$$\mathbb{E}S_a S_b(\mathbf{Zv} + \mathbf{e})_i = O(\sqrt{k/n}). \quad (3.24)$$

We now state our main theorem for this section.

Theorem 3.1 Under the Assumptions(1)-(4), if $d^2/n \rightarrow 0$, we have

$$\sup_{q \in \mathbb{R}} |\mathcal{L}_n(q) - \mathcal{L}_n^*(q)| = O_P(d^3 n^{-3/2}). \quad (3.25)$$

A direct application of Theorem 3.1 is the following result on highly accurate prediction intervals.

Theorem 3.2 Under the Assumptions(1)-(4) and $d^2/n \rightarrow 0$, for any $\alpha \in (0, 1)$, if q_1 and q_2 are real numbers such that

$$\begin{aligned} & \mathcal{L}_n^*(q_2) - \mathcal{L}_n^*(q_1) \\ & = 1 - \alpha, \end{aligned}$$

we have

$$\mathbb{P}[\hat{\mu}_T + q_1 \hat{\sigma}_T \leq T \leq \hat{\mu}_T + q_2 \hat{\sigma}_T] \quad (3.26)$$

$$= 1 - \alpha + O(d^3 n^{-3/2}). \quad (3.27)$$

Since the Fay-Herriot model (described in Section 2) is an important example, we state the results for it in a separate corollary below.

Corollary 3.1 In the Fay-Herriot model, assume that the matrix \mathbf{X} is full column rank, the diagonal entries h_{ii} of the projection matrix on the columns of \mathbf{X} satisfy $\sup_i h_{ii} = O(p/n)$, the Level 1 variances $\{D_i\}$ lie in a compact subset of $(0, \infty)$, and the estimator \hat{A} of A is positive. Then, for any $i \in \{1, \dots, n\}$, if $\hat{\theta}_i^{EB} = (1 - \hat{B}_i)Y_i + \hat{B}_i \mathbf{x}_i^T \hat{\beta}$, $\hat{\theta}_i^{EB*} = (1 - \hat{B}_i^*)Y_i + \hat{B}_i^* \mathbf{x}_i^T \hat{\beta}^*$, we have

$$\begin{aligned} \mathbb{P} \left[\theta_i \in \left(\hat{\theta}_i^{EB} + q_{i1} D_i^{1/2} (1 - \hat{B}_i)^{1/2}, \right. \right. \\ \left. \left. \hat{\theta}_i^{EB} + q_{i2} D_i^{1/2} (1 - \hat{B}_i)^{1/2} \right) \right] \\ = 1 - \alpha + O(p^3 n^{-3/2}); \end{aligned} \quad (3.28)$$

where $\hat{B}_i = D_i / (\hat{A} + D_i)$, and (q_{i1}, q_{i2}) satisfy

$$\begin{aligned} \mathbb{P}^* \left[\theta_i^* \in \left(\hat{\theta}_i^{EB*} + q_{i1} D_i^{1/2} (1 - \hat{B}_i^*)^{1/2}, \right. \right. \\ \left. \left. \hat{\theta}_i^{EB*} + q_{i2} D_i^{1/2} (1 - \hat{B}_i^*)^{1/2} \right) \right] \\ = 1 - \alpha + O_P(p^3 n^{-3/2}). \end{aligned}$$

The notations used in Corollary 3.1 are standard ones, that is, \mathbb{P}^* is the probability on the resampling scheme conditional on the data, $\hat{B}_i^* = D_i / (\hat{A}^* + D_i)$, where $\hat{\beta}^*$ and \hat{A}^* are the estimators computed on the bootstrap data \mathbf{Y}^* . Here conditional on the data, $\theta_i^* \sim N(\mathbf{x}_i^T \hat{\beta}, \hat{A})$, and $Y_i^* | \theta_i^* \sim N(\theta_i^*, D_i)$ independently. Corollary 3.1 is easily derived from Theorem 3.2. A slightly different approach to the same result may be found in the unpublished manuscript Chatterjee and Lahiri (2002). We now discuss the assumptions leading to our main result Theorem 3.1, and some additional features of our result.

REMARK 1. (On the dimension of the random effect vector:) Note that the dimension q of the random effect \mathbf{v} is arbitrary which may or may not depend on

n . Owing to this generalization, our analysis is for $T = c^T(\mathbf{X}\beta + \mathbf{Z}\mathbf{v})$, rather than the more traditional $\tilde{T} = c_1^T\beta + c_2^T\mathbf{v}$. Since \mathbf{X} is full column rank, the fixed effects in T and \tilde{T} are equivalent.

REMARK 2. (On the technical assumptions:) In the development of all the assumptions above, we have preferred simplicity over generality. The requirement $d^2n^{-1} \rightarrow 0$ is standard in dimension asymptotics. Assumption 1 ensures that both the fixed component and the variance of the random component of μ_T are $O(1)$, and the variance σ_T^2 is bounded away from zero and infinity. By suitably scaling the norm of the vector c this assumption is satisfied.

Assumption 2 is a standard assumption on the behavior \mathbf{X} . It ensures that the norm of each fixed effects covariate is of suitable order, and the fixed effects design is not singular. This assumption can be modified to suit cases where \mathbf{X} is not full column rank, but such generalizations are routine.

Assumption 3 is on standard differentiability and eigenvalue conditions. Here again, we have tried to adopt simple conditions rather than the most general ones. Note that the existence of the representations (3.13) and (3.14) are not part of the assumptions, and these representations will be established in the proof of Theorem 3.1.

Also note that the eigenvalues of $D(\hat{\psi})$ and $R(\hat{\psi})$ are estimates of the variance components in typical applications. Note that we do not allow these to be zero, since these must always lie in $(L^{-1}/2, 2L)$. However, L may be arbitrarily large, consequently this assumption does not limit the applicability of our results.

In Assumption 4 we take all moments of S to exist in order to achieve simplicity. Our result involves computation of several terms involving S , and having all the moments of S available simplifies the algebra. In most applications, both ψ and $\hat{\psi}$ lie in a compact set, hence this is not a strong condition. The other moment conditions on S given by (3.21)-(3.22) are routine. These hold when $\hat{\psi}$ is obtained using either maximum likelihood or restricted maximum likelihood formulation, see Jiang (1998) for related developments.

The conditions (3.23)-(3.24) are interesting, since they effectively set a limit to the amount of dependency structure we can have in Σ . In order to visualize this, suppose $\hat{\psi}^{(-i)}$ is the estimator of ψ obtained by using only those observations that are independent of Y_i ; and let $S^{(-i)} = (k/n)^{1/2}(\hat{\psi}^{(-i)} - \psi)$. Then, a sufficient condition for $\mathbb{E}S_j(\mathbf{Z}\mathbf{v} + \mathbf{e})_i = O((k/n)^{1/2})$ is that $S - S^{(-i)} = O_P((k/n)^{1/2})$.

This is routinely achieved, and in particular, if Y_i is independent of all but a finite number of observations, we have $S - S^{(-i)} = O_P((k/n)^{1/2})$. This is the typical situation is almost all applications of small area studies. Thus, the effect of Assumption 4 is to restrict the complexity of the matrices D and R .

REMARK 3. (On the nature of prediction intervals:) A prominent application of the highly accurate approximation of $\mathcal{L}_n(\cdot)$ by $\mathcal{L}_n^*(\cdot)$ is stated in Theorem 3.2,

that is, in the construction of prediction intervals. Note that these are *bootstrap intervals*, as opposed to the traditional intervals described in Section 2, some of which are improved with bootstrap corrections.

However, Theorem 3.2 does not describe the nature of the bootstrap prediction intervals, since the choice of q_1 and q_2 can be quite arbitrary. These may be chosen to ensure either an *equal tail* property of the interval; whereby $\mathcal{L}_n^*(q_1) = \alpha/2$ and $\mathcal{L}_n^*(q_2) = 1 - \alpha/2$; or we may chose these according to a *minimum length of interval* property, that is, we minimize the length $\hat{\sigma}_T(q_2 - q_1)$. The simulation experiments reported in Section 4 show that both equal tailed and minimum length bootstrap prediction intervals typically achieve the desired coverage accuracy without the use of elaborate calibrations; and the minimum length interval is always slightly shorter than the equal tailed one.

REMARK 4. (On multivariate prediction:) Note that in place of the real valued T studied above, we could have a vector valued T with little change in methodology. The algorithmic and algebraic details are similar, and the main result of high order accuracy of distributional approximation (3.25) holds. The major difference between univariate and multivariate prediction is in the construction of prediction regions. Instead of the two points q_1 and q_2 , we need to obtain probability concentration regions from the bootstrap distribution. Such regions can be obtained using various data depth notions and shape features, for example, as in Yeh and Singh (1997). This is a separate issue from the one addressed in this paper, and will be handled in a different paper. Note that multidimensional probability concentration regions can be quite hard to calibrate in practice. Some techniques, such as calibration of the end points of an interval, are not available in this case.

REMARK 5. (Asymptotics on total sample size n :) One important feature of Theorem 3.1 is that the asymptotic limits are obtained with total sample size n tending to infinity. The total sample size n is the sum total of all observations made, counting each repeated measurement on each individual unit in each small area as a distinct observation. This allows Theorem 3.1 to be used with considerable flexibility, for example, when number of individual units in small areas are large, or when number of small areas are large, or both. However, requirements of asymptotic negligibility, as in (3.23)-(3.24), must still be met. Our assumptions are designed for the more realistic applications where number of small areas are large.

In general, for linear mixed models asymptotic limits are obtained either when the number of observations in each small area tends to infinity, or when the number of small areas tend to infinity, see McCulloch and Searle (2001) and Rao (2003) for details. Theorem 3.1 is a breakthrough, owing to the greater flexibility it allows in asymptotics.

REMARK 6. (On area specific properties:) The *area-specific signal* for each small area is $T_i = \mathbf{x}_i^T\beta + \mathbf{Z}_i^T\mathbf{v}$, conditional on the observed i^{th} small area data Y_i . Dis-

tributions of predictors for such area specific signals are effectively captured by our bootstrap predictive distribution approximation. Consequently, our bootstrap prediction intervals for T_i are also *area specific*. Extensions to compare two or more small areas can be obtained by similar techniques, see comment on multidimensional prediction above.

In the prediction interval described in Theorem 3.2, we have considered unconditional coverage, where probabilities are computed over the joint distribution of \mathbf{Y} and \mathbf{v} . This establishes the performance of the area-specific interval that depends on \mathbf{v} conditional on \mathbf{Y} , as well as variability due to observations \mathbf{Y} .

Alternatively, one might compute the *area-specific (random) coverage*, which is defined as $\mathbb{P}[T_i \in I_{P_i} | \mathbf{Y}_i]$, where I_{P_i} is the prediction interval. The interval proposed in Theorem 3.2 achieves $O_P(d^2/n)$ order of area-specific coverage accuracy, since some smoothing effects arising from the distribution of \mathbf{Y} are absent. This is no worse (and in some cases, better) than the area-specific coverage obtained by other techniques in special cases of the general linear model (3.4).

REMARK 7. (On calibration:) Both the unconditional as well as the area-specific coverage can be improved by calibration. The use of calibration coupled with resampling is an active topic of research, and some discussion on this has been presented in Section 2. The coverage accuracy of the prediction interval of Theorem 3.2 can be improved to $O(d^5 n^{-5/2})$ with one round of calibration, and further still with more calibration. Such calibration may be done either on the probabilities corresponding to the two end points as in DiCiccio and Efron (1996), or on the true coverage of the interval. Some of our simulations, not reported in this paper, suggest that it is not always beneficial to attempt boosting the theoretical coverage probability, disregarding other properties of the interval. For example, variability of calibrated intervals are greater than uncalibrated ones, minimum length property is almost never preserved, and the results are quite dependent on the parameters and fixed constants of the problem. Hence, it seems reasonable to work with a good predictive distribution as in Theorem 3.1, instead of starting with a naive interval and embarking on intense iterative calibration.

REMARK 8. (The parallel work of Hall and Maiti (2006b):) Recently, Hall and Maiti (2006b) studied parametric bootstrap methods for general small area models, and considering the overlap of the topics studied in their paper and this one, deserve special mention. For this comment, we use some notations from Hall and Maiti (2006b) whenever they are not in conflict with the notations in the rest of this paper, but use our notations otherwise.

For a suitable function $f_i(\beta)$ involving co-variates $X_i = (X_{i1}, \dots, X_{in_i})$ and parameter β , they consider random effect $\Theta_i \sim Q(\cdot; f_i(\beta), \xi)$, and conditional on Θ_i , the data Y_{ij} are independent observations from $R(\cdot; \psi(\Theta_i), \eta_i)$, for $j = 1, \dots, n_i$, $i = 1, \dots, m$. Here $\psi(\cdot)$ is a known link

function, ξ and η_i 's are either parameters or known constants, and $Q(\cdot)$ and $R(\cdot)$ are known probability distribution functions. They go on to study calibration of the mean squared prediction error (MSPE) and interval estimation with parametric bootstrap.

Their model is broad enough to handle nonlinear mixed effects, which our model (1.1) cannot do. However, their assumption of Y_{ij} 's being independent means that they do not consider longitudinal models, or other models with temporal or spatial dependence. This is essentially the case R being a multiple of the identity matrix in our set-up. Our model is broader than Hall and Maiti's in including several varieties of dependence structure.

The interval estimate from their Section 2.8 is

$$\hat{I}_\alpha = \mathbf{x}_i^T \hat{\beta} \pm z_{\alpha/2} \hat{A}^{1/2} \quad (3.29)$$

for the Fay-Herriot model. Rao (2005) noted that this interval does not make use of the area-specific direct estimator, unlike the prediction interval proposed by Chatterjee and Lahiri (2002). Hall and Maiti (2006b) calibrate this interval for better coverage accuracy, improving from their result

$$\mathbb{P}[\Theta_i \in \hat{I}_\alpha] = 1 - \alpha + O(m^{-1}). \quad (3.30)$$

The result (3.30) hold when the probability statement is on the marginal distribution of random effect Θ_i , and estimators $\hat{\beta}$ and \hat{A} are independent of i^{th} area data Y_{i1}, \dots, Y_{in_i} .

Our probability statements in Theorem 3.1 and Theorem 3.2 are, however, on the *joint variability* of the random effects and data $(\Theta_i, Y_{i1}, \dots, Y_{in_i})$. Also note that Theorem 3.2 is obtained as $n = \sum_{i=1}^m n_i \rightarrow \infty$, while (3.30) is obtained as $m \rightarrow \infty$. Since some of the n_i 's can be large, the speed of convergence towards the asymptotic limits are different; and $m = o(n)$ if any $n_i \rightarrow \infty$.

Hall and Maiti (2006b) obtain that if \hat{I}_α is calibrated once (twice), the coverage accuracy improves to $O(m^{-2})$ ($O(m^{-3})$). If the interval in (3.27) or (3.28) is calibrated once (twice), the coverage accuracy improves to $O(n^{-5/2})$ ($O(n^{-7/2})$) when parameter dimension is fixed.

In summary, Hall and Maiti (2006b) cover a wide ranging independent data framework, with careful MSPE estimation and marginal coverage of prediction intervals as number of small areas increases; while we consider deeper linear mixed framework allowing for longitudinal dependence, and establish results as total data size increases on the joint variability of random effects and data, thus also obtaining area specificity.

4 A simulation example

In this section we compare the performance of our proposed parametric bootstrap with that of the traditional approaches, using a simulation study. For the sake of comparability with existing studies, we adopt part of the simulation framework of Datta, Rao, Smith (2005) for our study. We consider the Fay-Herriot model with $m = 15$

and $\mathbf{x}_i^T \beta = 0$. But, to account for the uncertainty in estimation of the common mean that arises in practice, we still estimate the zero mean.

We consider five groups G_t , ($t = 1, \dots, 5$) of small areas with three areas in each group. Within each group, sampling variances are the same. Specifically, $D_j = c_1$ $j \in G_1$; $D_j = c_2$ $j \in G_2$; $D_j = c_3$ $j \in G_3$; $D_j = c_4$ $j \in G_4$; $D_j = c_5$ $j \in G_5$. Let $c = (c_1, \dots, c_5)'$. We consider the following four different patterns of A and D_i 's combinations: (a) $A = 1; c = (4.0, 0.6, 0.5, 0.4, 0.2)'$ [this is pattern (c) of Datta et al. (2005)]; (b) $A = 2; c = (4.0, 0.6, 0.5, 0.4, 0.2)'$; (c) $A = 1; c = (8.0, 1.2, 1.0, 0.8, 0.4)'$, and (d) $A = 2; c = (8.0, 1.2, 1.0, 0.8, 0.4)'$.

We obtain all the results based on 10,000 simulation runs. The Prasad-Rao method-of-moments, and the Fay-Herriot method of estimating the variance component A are considered. Tables 4.1-4.4 report the simulated coverage probabilities and average lengths of several different prediction intervals (with nominal coverage 0.95). We consider three prediction intervals of the type $EBLUP \pm 1.96\sqrt{mspe}$, where $mspe$ is an estimator of the MSPE of EBLUP. The Cox interval, discussed in Section 2, is obtained by using the Prasad-Rao method-of-moment estimator of A . The Prasad-Rao (PR) interval estimator is obtained using that estimator of A along with the Prasad-Rao (1990) MSPE estimator. The Fay-Herriot (FH) interval estimator is obtained by using Fay-Herriot method of moments estimator of A [see Fay and Herriot 1979], and the MSPE estimator of EBLUP considered by Datta, Rao and Smith (2005). Along with these three, we report two different parametric bootstrap prediction intervals. In both the methods, we used Fay-Herriot method of estimating A , and the ordinary least squares estimator of β . Whenever the estimate of A is negative, we truncate it at $1/6$, using a truncation rule suggested by Morris (1983a). The first bootstrap interval is equal-tailed (PB-ET), and the second is the shortest length prediction interval (PB-SL). For both cases, we considered bootstrap sample of size 1000.

The figures in Table 4.1-4.4 are the average coverage probabilities and average lengths for each prediction interval method for patterns (a)-(d), the average being taken over all three small areas within each group. It is clear that the results depend on the patterns of A and D_i 's. The Cox prediction interval method consistently undercovers. The Fay-Herriot method interval improves on the Cox interval, but is still subject to the undercoverage problem. In general, the Prasad-Rao method interval is overly conservative at the expense of large average length. The extent of the undercoverage for the Fay-Herriot interval or the overcoverage for the Prasad-Rao method depends very much on the pattern of A and D_i 's. It appears that the problem is more severe for group G_1 where the ratio D_i/A is the maximum. Thus, while both the Prasad-Rao and the Fay-Herriot MSPE estimators enjoy good theoretical properties, the resulting interval estimates suffer from coverage problem owing to the en-

forced symmetry and normality assumption. In contrast, the performances of our parametric bootstrap methods remain stable over all four different patterns and always close to the target nominal level. Our minimum length parametric bootstrap method tends to provide shorter prediction intervals compared to the equal-tailed equivalents.

References

- [1] ABRAMOVITCH, L. AND SINGH, K., (1985) Edgeworth corrected pivotal statistics and the bootstrap, *Ann. Statist.*, **13**, 116-132.
- [2] AITCHISON, J. (1975) Goodness of predictive fit, *Biometrika*, **62**, 547-554.
- [3] BASU, R., GHOSH, J. K., AND MUKERJEE, R., (2003), Empirical Bayes prediction intervals in a normal regression model: higher order asymptotics, *Statist. Probab. Lett.*, **63**, 197-203.
- [4] BERAN, R. (1990a) Refining bootstrap simultaneous confidence sets. *J. Amer. Statist. Assoc.* **85**, 417-426.
- [5] BERAN, R. (1990b) Calibrating prediction regions. Refining bootstrap simultaneous confidence sets. *J. Amer. Statist. Assoc.* **85**, 715-723.
- [6] BREIMAN, L. (1996) Bagging predictors, *Machine Learning*, **24**, 123-140.
- [7] CARLIN, B.P. AND LOUIS, T.A. (1996), *Bayes and empirical Bayes methods for data analysis*, Chapman and Hall, London.
- [8] CARLIN, B. AND GELFAND, A. (1990), Approaches for empirical Bayes confidence intervals, *J. Amer. Statist. Assoc.*, **85**, 105-114.
- [9] CARLIN, B. AND GELFAND, A. (1991), A sample reuse method for accurate parametric empirical Bayes confidence intervals, *J. Roy. Statist. Soc. Ser. B*, **53**, 189-200.
- [10] CHATTERJEE, S. AND LAHIRI, P., (2002), Parametric bootstrap confidence intervals in small area estimation problems, *unpublished manuscript*.
- [11] Chatterjee, A. Lahiri, P., and Li, H. (2007), "Parametric bootstrap approximation to the distribution of EBLUP, and related prediction intervals in linear mixed models," *Annals of Statistics*, in press.
- [12] COX, D.R. (1975), Prediction intervals and empirical Bayes confidence intervals, in *Perspectives in Probability and Statistics*, papers in honor of M.S. Bartlett. Ed: J. Gani, Academic Press, 47-55.
- [13] DAS, K., JIANG, J., AND RAO, J. N. K., (2004), Mean squared error of empirical predictor, *Ann. Statist.*, **32**, 818-840.

- [14] DATTA, G.S., GHOSH, M., SMITH, D., AND LAHIRI, P. (2002), On an asymptotic theory of conditional and unconditional coverage probabilities of empirical Bayes confidence intervals, *Scand. J. Statist.* **29**, 139-152.
- [15] DATTA, G.S., RAO, J.N.K. AND SMITH, D.D. (2005), On measuring the variability of small area estimators under a basic area level model, *Biometrika*, **92**, 183-196.
- [16] DICICCIO, T. AND EFRON, B. Bootstrap confidence intervals (with discussion), *Statist. Sci.* **11**, No. 3, 189-228.
- [17] EFRON, B. AND TIBSHIRANI, R. J. (1993) *An introduction to the bootstrap*, Chapman and Hall, New York.
- [18] FAY, R. E. AND HERRIOT, R. A. (1979) Estimates of income for small places: an application of James-Stein procedure to census data, *J. Amer. Statist. Assoc.*, **74**, 269-277.
- [19] FUSHIKI, T., KOMAKI, F. AND AIHARA, K. (2004) On parametric bootstrapping and Bayesian prediction, *Scand. J. Statist.*, **31**, 403-416.
- [20] FUSHIKI, T., KOMAKI, F. AND AIHARA, K. (2005) *Bernoulli*, **11**, 293-307.
- [21] GEORGE, E. I., LIANG, F. AND XU, X. (2006) Improved minimax predictive densities under Kullback-Leibler loss, *Ann. Statist.*, **34**, 78-91.
- [22] HALL, P. (2006) Discussion of *Mixed model prediction and small area estimation*, *Test*, **15**, 1-96.
- [23] HALL, P. AND MAITI, T. (2006a). Nonparametric estimation of mean-squared prediction error in nested-error regression models. *Ann. Statist.*, **34**, 1733-1750
- [24] HALL, P. AND MAITI, T. (2006b). On parametric bootstrap methods for small-area prediction. *J. Roy. Statist. Soc. Ser. B*, **68**, 221-238.
- [25] HALL, P. AND MARTIN, M. A. (1996). Discussion on *Bootstrap confidence intervals* by DiCiccio and Efron, *Statist. Sci.*, **11**, No. 3, 212-214.
- [26] HARRIS, I. R. (1989) Predictive fit for natural exponential families, *Biometrika*, **76**, 675-684.
- [27] HARTIGAN, J. (1964) Invariant prior distributions, *Ann. Math. Statist.*, **35**, 836-845.
- [28] HARTIGAN, J. (1998) The maximum likelihood prior, *Ann. Statist.*, **26**, 2083-2103.
- [29] HILL, J.R. (1990), A general framework for model-based statistics, *Biometrika*, **77**, 115-126.
- [30] JESKE, D. R. AND HARVILLE, D. A. (1988), Prediction-interval procedures and (fixed-effects) confidence-interval procedures for mixed linear models, *Commun. Statist. - Theory Meth.* **17**, 1053-1087.
- [31] Jiang, J. (1996), "REML estimation: asymptotic behavior and related topics," *Annals Statistics* **24**, 255-286.
- [32] Jiang, J. (1998), "Asymptotic properties of the empirical BLUP and BLUE in mixed linear models," *Statistica Sinica* **8**, 861-885.
- [33] JIANG, J., LAHIRI, P. AND WAN, S. (2002), A unified jackknife theory for empirical best prediction with M-estimation, *Ann. Statist.* **30**, 1782-1810.
- [34] JIANG, J., AND LAHIRI, P. (2006). Mixed model prediction and small area estimation (with discussions), *Test*, **15**, 1, 1-96.
- [35] JIANG, J. AND ZHANG, W. (2002), Distribution-free prediction intervals in mixed linear models, *Statistica Sinica* **12**, 537-553.
- [36] KOMAKI, F. (1996) On asymptotic properties of predictive distributions, *Biometrika*, **83**, 299-313.
- [37] KOMAKI, F. (2001) A shrinkage predictive distribution for multivariate normal observations, *Biometrika*, **88**, 859-864.
- [38] KOMAKI, F. (2006) Shrinkage priors for Bayesian prediction, *Ann. Statist.*, **34**, 808-819.
- [39] LAIRD, N. M. AND LOUIS, T. A. (1987) Empirical Bayes confidence intervals based on bootstrap samples (with discussion), *J. Amer. Statist. Assoc.*, **82**, 739-750.
- [40] LEE, S. M S. AND YOUNG, G. A. (1996). Discussion on *Bootstrap confidence intervals* by DiCiccio and Efron, *Statist. Sci.*, **11**, No. 3, 221-223.
- [41] MCCULLOCH, C.E. AND SEARLE, S. R., (2001), *Generalized, Linear, and Mixed Models*, Wiley, New York.
- [42] MORRIS, C.N., (1983a) Parametric empirical Bayes inference: theory and applications (with discussion), *J. Amer. Statist. Assoc.*, **78**, 47-65.
- [43] MORRIS, C.N., (1983b) Parametric empirical Bayes confidence intervals, in *Scientific Inference, Data Analysis and Robustness*, Academic Press, New York, 25-50.
- [44] MURRAY, D. G. (1977) A note on the estimation of probability density functions, *Biometrika*, **64**, 150-152.
- [45] NG, V. M. (1980) On the estimation of parametric density functions, *Biometrika*, **67**, 505-506.

- [46] PRASAD, N. G. N. AND RAO, J. N. K., (1990), The estimation of the mean squared error of small-area estimators, *J. Amer. Statist. Assoc.*, **85**, 163-171.
- [47] RAO, C. R., (1963), *Linear statistical inference and its applications*, Wiley Eastern, New Delhi.
- [48] RAO, J. N. K., (2003), *Small Area Estimation*, Wiley, New York.
- [49] RAO, J.N.K., (2005), Inferential issues in small area estimation: some new developments, *Statistics in Transition*, Vol 7, 2003, 513-526.
- [50] YEH, A. B. AND SINGH, K. (1997) Balanced confidence regions based on Tukey's depth and the bootstrap. *J. Roy. Statist. Soc. Ser. B*, **59**, 639-652.

Table 4.1: $A = 1$, $D = (4.0, 0.6, 0.5, 0.4, 0.2)$, $N = 10,000$, $BN = 1000$

Group	Cox	FH	PR	PB-ET	PB-SL
G1	83.1 (3.12)	90.4 (3.57)	92.4 (3.82)	96.1 (4.50)	95.7 (4.42)
G2	85.4 (2.14)	93.7 (2.50)	98.0 (3.19)	96.2 (2.83)	95.9 (2.79)
G3	85.8 (2.02)	93.9 (2.36)	98.0 (3.08)	96.0 (2.65)	95.6 (2.61)
G4	86.1 (1.89)	94.3 (2.19)	98.2 (2.93)	96.1 (2.43)	95.7 (2.39)
G5	89.7 (1.12)	95.2 (1.23)	97.3 (1.87)	95.7 (1.28)	95.3 (1.26)

Table 4.2: $A = 2$, $D = (4.0, 0.6, 0.5, 0.4, 0.2)$, $N = 10,000$, $BN = 1000$

Group	Cox	FH	PR	PB-ET	PB-SL
G1	88.0 (4.18)	92.7 (4.64)	92.7 (4.73)	96.2 (5.29)	95.9 (5.21)
G2	90.6 (2.51)	94.7 (2.74)	96.1 (2.97)	96.0 (2.92)	95.7 (2.88)
G3	90.7 (2.35)	94.7 (2.55)	96.1 (2.78)	95.9 (2.70)	95.5 (2.67)
G4	91.1 (2.16)	94.9 (2.32)	96.3 (2.55)	96.0 (2.45)	95.8 (2.42)
G5	92.7 (1.18)	94.8 (1.23)	95.6 (1.40)	95.5 (1.26)	95.0 (1.25)

Table 4.3: $A = 1$, $D = (8.0, 1.2, 1.0, 0.8, 0.4)$, $N = 10,000$, $BN = 1000$

Group	Cox	FH	PR	PB-ET	PB-SL
G1	78.0 (3.28)	86.7 (3.72)	93.9 (4.51)	94.5 (5.34)	94.3 (5.24)
G2	79.2 (2.45)	91.9 (3.06)	99.3 (5.03)	95.2 (3.84)	94.9 (3.77)
G3	79.5 (2.34)	92.3 (2.95)	99.2 (5.04)	95.1 (3.64)	94.8 (3.58)
G4	80.2 (2.21)	93.4 (2.81)	99.3 (5.03)	95.3 (3.39)	94.9 (3.33)
G5	84.4 (1.43)	96.2 (1.80)	98.5 (4.26)	95.6 (1.93)	95.3 (1.90)

Table 4.4: $A = 2$, $D = (8.0, 1.2, 1.0, 0.8, 0.4)$, $N = 10,000$, $BN = 1000$

Group	Cox	FH	PR	PB-ET	PB-SL
G1	81.4 (4.34)	90.3 (5.04)	91.7 (5.35)	96.5 (6.46)	96.2 (6.35)
G2	83.4 (2.96)	93.6 (3.53)	98.1 (4.58)	96.7 (4.11)	96.3 (4.05)
G3	83.8 (2.81)	93.8 (3.33)	98.1 (4.44)	96.5 (3.85)	96.2 (3.79)
G4	84.1 (2.62)	94.3 (3.09)	98.2 (4.27)	96.7 (3.53)	96.4 (3.47)
G5	87.6 (1.56)	95.2 (1.74)	97.3 (2.97)	96.1 (1.87)	95.7 (1.85)