

University of Wollongong

Research Online

Faculty of Informatics - Papers (Archive)

Faculty of Engineering and Information
Sciences

1-1-2008

Adaptive Calibration for Prediction of Finite Population Totals

Raymond L. Chambers

University of Wollongong, ray@uow.edu.au

Robert Graham Clark

University of Wollongong, rclark@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>



Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Chambers, Raymond L. and Clark, Robert Graham: Adaptive Calibration for Prediction of Finite Population Totals 2008, 163-172.

<https://ro.uow.edu.au/infopapers/1404>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Adaptive Calibration for Prediction of Finite Population Totals

Abstract

Sample weights can be calibrated to reflect the known population totals of a set of auxiliary variables. Predictors of finite population totals calculated using these weights have low bias if these variables are related to the variable of interest, but can have high variance if too many auxiliary variables are used. This article develops an adaptive calibration approach, where the auxiliary variables to be used in weighting are selected using sample data. Adaptively calibrated estimators are shown to have lower mean squared error and better coverage properties than non-adaptive estimators in many cases.

Keywords

sample surveys, sample weighting; prediction approach; ridge estimation; model selection; stepwise procedures

Disciplines

Physical Sciences and Mathematics

Publication Details

Clark, R. G. & Chambers, R. L. (2008). Adaptive Calibration for Prediction of Finite Population Totals. *Survey Methodology*, 34 (2), 163-172.

Article

Adaptive calibration for prediction of finite population totals

by Robert G. Clark and Raymond L. Chambers

December 2008



Adaptive calibration for prediction of finite population totals

Robert G. Clark and Raymond L. Chambers¹

Abstract

Sample weights can be calibrated to reflect the known population totals of a set of auxiliary variables. Predictors of finite population totals calculated using these weights have low bias if these variables are related to the variable of interest, but can have high variance if too many auxiliary variables are used. This article develops an “adaptive calibration” approach, where the auxiliary variables to be used in weighting are selected using sample data. Adaptively calibrated estimators are shown to have lower mean squared error and better coverage properties than non-adaptive estimators in many cases.

Key Words: Sample surveys; Sample weighting; Prediction approach; Ridge estimation; Model selection; Stepwise procedures.

1. Introduction

Predictors of finite population totals are commonly calculated by weighted sums of sample values. Auxiliary variables are often available, whose sample values and population totals are known. Weights can be constructed so that weighted sums of auxiliary variables agree with the known population totals, a process called calibration (Deville and Särndal 1992). Predictors of finite population totals based on calibrated weights generally have much lower prediction bias than predictors calculated without auxiliary information.

Existing literature on finite population prediction essentially assumes that a set of useful auxiliary variables is chosen without reference to sample data. In practice, however, there may be a large set of potential auxiliary variables, not all of which should be used. Using additional auxiliary variables generally reduces the bias of calibrated predictors but increases the variance, so that using too many auxiliary variables can actually increase the mean squared error of calibrated predictors. The choice of which auxiliary variables to use is often not obvious, and sample data may be required to determine which set of auxiliary variables is appropriate for predictors of the totals of particular variables of interest. This paper develops methods for making this determination. Our approach may be called adaptive calibration, because the set of variables is chosen adaptively from sample data, rather than statically without reference to the sample at hand.

The prediction framework to finite population estimation will be used (see for example Brewer 1963; Royall 1970; Valliant, Dorfman and Royall 2000). In this approach, the population values of the variables of interest are treated as random variables. The aim is to predict the population total (which is also a random variable) or other finite population quantities using sample data on the variable of interest, and

population data on some auxiliary variables. The sample may have been selected using probability sampling or some other method, and is conditioned upon in inference. A stochastic model for the variable of interest is a central feature. One feature of the prediction framework is that misspecification of the model, for example due to omitting important auxiliary variables, can lead to substantial bias.

An alternative framework is the model-assisted approach (Särndal, Swensson and Wretman 1992). In this approach, a stochastic model is used but the model plays a less crucial role. The randomized nature of sampling is exploited to ensure that estimators are approximately unbiased even if the model is incorrect. When the model is correct, both approaches give approximately unbiased estimators, but the model-based approach would generally give lower variances of estimators of interest. If the model is misspecified, then model-based predictors and variance estimators may be more biased, however robust model-based methods have been developed to combat this problem. For example Royall and Herson (1973a, 1973b) discuss robust prediction and Royall and Cumberland (1978, 1981a, 1981b) developed variance estimators that are robust to heteroscedasticity. For comparisons of the prediction and model-assisted frameworks, see for example Smith (1976) and Hansen, Madow and Tepping (1983).

The problem of selecting a set of auxiliary variables in the model-assisted framework was considered by Silva and Skinner (1997) and Skinner and Silva (1997). They found that adding calibration variables reduces the mean squared error (MSE) up to a point, after which adding further variables increases the MSE. Choosing calibration variables adaptively, based on sample data, gave better estimates than either calibrating on all variables or no variables. The applicability of this work to model-based prediction is not clear, because the role of the model is very different in the two frameworks. Misspecified models can lead to

1. Robert G. Clark and Raymond L. Chambers, Centre for Statistical and Survey Methodology, University of Wollongong, NSW 2522 Australia. E-mail: Robert_Clark@uow.edu.au.

substantially biased model-based predictors, whereas model-assisted estimators are approximately unbiased even if important variables are omitted. As a result, different strategies for model selection could be appropriate in the two frameworks. Moreover, the differences between alternative approaches would be expected to be more pronounced in the prediction framework than in the model-assisted framework.

Chambers, Skinner and Wang (1999) proposed an approach for selecting calibration variables in the prediction framework, using forward, backward or stepwise selection. (This paper will henceforth be referred to as CSW.) The decision whether to omit (or add) a variable at each step was based on minimizing the estimated squared error of prediction (MSEP) for the predictor of interest. The approach was not evaluated by simulation study, and the estimators of MSEP used were not robust to heteroscedasticity.

The purpose of this paper is to develop the basic approach of CSW to apply to a wider range of situations, including heteroscedastic populations and multi-stage samples, and to evaluate the approach using realistic simulation studies. Estimators of the MSEP which are robust to heteroscedasticity, and to correlation in the case of multi-stage surveys, will be used. The performance of the estimators will be evaluated by simulation from two populations: financial data on farms generated from a farm survey and labour force data from a population census.

Following CSW, the basic approach will be to build a set of auxiliary variables using stepwise selection of variables, starting with some initial set. This algorithm builds up a set of auxiliary variables by a sequence of many decisions between two nested sets of variables. We compare several alternative criteria for deciding between two nested sets, including statistical significance and a number of alternative estimators of the mean squared error of prediction (MSEP). Three alternative estimators of MSEP are considered: a non-robust estimator; an estimator of MSEP which is robust to heteroscedasticity; and an estimator which is robust both to heteroscedasticity and correlations within primary sampling units in multi-stage sampling.

Section 2 contains notation and definitions. Section 3 derives the difference in the MSEP of two predictors based on nested models, and develops several alternative estimators of this difference. Section 4 contains simulation results for a farm survey and a multi-stage household survey. Section 5 is a discussion. We conclude that adaptive calibration generally performs better than static calibration, provided that a non-robust estimator of the MSEP, or statistical significance, is used as the objective in model selection.

2. Notation and definitions

A variable of interest Y_i is observed for a sample s of n units, which is a subset of a finite population U containing N units. The aim is to estimate the population total $T_Y = \sum_{i \in U} Y_i$ and other finite population quantities of Y . A vector of auxiliary variables \mathbf{x}_i is available for $i = 1, \dots, n$, with known population total $\mathbf{T}_x = \sum_{i \in U} \mathbf{x}_i$.

Weighted estimators of T_Y are given by $\hat{T}_Y = \sum_{i \in s} w_i Y_i$, where w_i can depend on the auxiliary variables but not on the variable of interest. A set of weights is said to be calibrated on \mathbf{x}_i if $\sum_{i \in s} w_i \mathbf{x}_i = \mathbf{T}_x$.

The best linear unbiased predictor (BLUP) based on a linear regression model is one example of a calibrated estimator. The most commonly used BLUP is based on the model

$$\begin{aligned} E[Y_i] &= \boldsymbol{\beta}^T \mathbf{x}_i \\ \text{var}[Y_i] &= \sigma_i^2 = v_i \sigma^2 \end{aligned} \quad (1)$$

$$\text{cov}[Y_i, Y_j] = 0 (i \neq j)$$

(with v_i assumed to be known) and is given by

$$\hat{T}_Y = \sum_{i \in s} Y_i + \sum_{i \in r} \hat{\boldsymbol{\beta}}^T \mathbf{x}_i \quad (2)$$

where $r = U - s$ is the set of non-sample units and

$$\hat{\boldsymbol{\beta}} = \left\{ \sum_{i \in s} v_i^{-1} \mathbf{x}_i \mathbf{x}_i^T \right\}^{-1} \sum_{i \in s} v_i^{-1} \mathbf{x}_i Y_i \quad (3)$$

is a weighted least squares estimator of $\boldsymbol{\beta}$. The BLUP can also be written in weighted form as

$$\hat{T}_Y = \sum_{i \in s} w_i Y_i$$

where the weights w_i are given by

$$w_i = 1 + \mathbf{T}_{xr}^T \left\{ \sum_{j \in s} v_j^{-1} \mathbf{x}_j \mathbf{x}_j^T \right\}^{-1} v_i^{-1} \mathbf{x}_i \quad (4)$$

and $\mathbf{T}_{xr} = \sum_{i \in r} \mathbf{x}_i$. It is straightforward to verify that $\sum_{i \in s} w_i \mathbf{x}_i = \mathbf{T}_x$.

For heteroscedastic data, it is usually difficult to model v_i reliably. In this case, robust estimators of the prediction variance of the BLUP are available, which do not rely on knowledge of v_i (Royall and Cumberland 1978, 1981a, 1981b). For multi-stage samples, the assumption of independence may be violated. In this case, the BLUP based on (1) may still be used, and a robust ultimate cluster variance estimator of its prediction variance can be used (e.g., Valliant *et al.* 2000, Chapter 9). An alternative approach, which will not be considered here, would be to construct a BLUP based on a model that includes the

within-cluster correlations (Royall 1976). Section 3 will discuss robust and non-robust estimation of the mean squared error of prediction of the BLUP in more detail.

A decision needs to be made on what to include in \mathbf{x}_i in the BLUP. Stepwise selection, forward selection and backward selection are algorithms that can be used to decide which subset of the available auxiliary variables should be used. All three algorithms include many choices between two nested sets of auxiliary variables. Suppose the choice is between (A) using a predictor \hat{T}_A based on \mathbf{x}_i and (B) using a predictor \hat{T}_B based on a subvector \mathbf{x}_{1i} . We can partition \mathbf{x}_i as $\mathbf{x}_i = (\mathbf{x}_{1i}^T, \mathbf{x}_{2i}^T)^T$. The number of elements of \mathbf{x}_i , \mathbf{x}_{1i} and \mathbf{x}_{2i} are denoted by p , p_1 and p_2 , respectively.

We similarly partition $\boldsymbol{\beta}$ as $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T$. Predictor \hat{T}_A is unbiased under model A:

$$E[Y_i] = \boldsymbol{\beta}^T \mathbf{x}_i = \boldsymbol{\beta}_1^T \mathbf{x}_{1i} + \boldsymbol{\beta}_2^T \mathbf{x}_{2i}. \quad (5)$$

The predictor \hat{T}_B is unbiased for model B,

$$E[Y_i] = \boldsymbol{\beta}_1^T \mathbf{x}_{1i}, \quad (6)$$

which is the special case of model A where $\boldsymbol{\beta}_2 = \mathbf{0}$.

3. Estimation of the difference in the MSEP

3.1 Comparing predictors from nested models

Following CSW, our approach is to estimate the difference in the MSEPs of the two estimators:

$$\Delta = E[(\hat{T}_A - T_Y)^2] - E[(\hat{T}_B - T_Y)^2]$$

where the expectations are evaluated with respect to model A, because model B is a special case of this model. Typically, \hat{T}_A will be less biased than \hat{T}_B but have higher variance. Either predictor can have higher or lower MSEP depending on the particular population and sample.

For single stage sampling, it is usually reasonable to assume Y_i and Y_j independent for all $i \neq j$. Section 3.2 will derive Δ and an estimator of it in this case. Section 3.3 will describe the instructive special case where variances are equal and BLUPs are used; this was the case considered by CSW. Section 3.4 extends this by describing a heteroscedasticity-robust estimator of Δ . Section 3.5 further extends the approach by deriving Δ and an estimator of it for multi-stage sampling where there may be correlations between values from the same cluster.

3.2 Estimating Δ in single-stage sampling with known variance

In addition to model (5), we assume in this subsection that Y_i and Y_j are independent for $i \neq j$ and that

$\text{var}[Y_i] = \sigma_i^2 = \sigma^2 v_i$ where v_i are known. In this case, the MSEP of any predictor $\hat{T} = \sum_{i \in S} w_i Y_i$ is given by

$$\begin{aligned} \text{MSEP}[\hat{T}] &= E[(\hat{T} - T_Y)^2] \\ &= \left\{ E \left[\sum_{i \in S} w_i Y_i - \sum_{i \in U} Y_i \right] \right\}^2 + \text{var} \left[\sum_{i \in S} (w_i - 1) Y_i - \sum_{i \in r} Y_i \right] \\ &= \left\{ \boldsymbol{\beta}^T \left(\sum_{i \in S} w_i \mathbf{x}_i - \sum_{i \in U} \mathbf{x}_i \right) \right\}^2 + \sum_{i \in S} (w_i - 1)^2 \sigma_i^2 + \sum_{i \in r} \sigma_i^2. \end{aligned}$$

Writing $\mathbf{d} = \sum_{i \in S} w_i \mathbf{x}_i - \mathbf{T}_x$, we can rewrite the MSEP as

$$\text{MSEP}[\hat{T}] = \mathbf{d}^T (\boldsymbol{\beta} \boldsymbol{\beta}^T) \mathbf{d} + \sum_{i \in S} (w_i - 1)^2 \sigma_i^2 + \sum_{i \in r} \sigma_i^2.$$

Let $\mathbf{d}_A = \sum_{i \in S} w_{Ai} \mathbf{x}_i - \mathbf{T}_x$ and $\mathbf{d}_B = \sum_{i \in S} w_{Bi} \mathbf{x}_i - \mathbf{T}_x$. Then Δ is given by:

$$\begin{aligned} \Delta &= \text{MSEP}[\hat{T}_A] - \text{MSEP}[\hat{T}_B] \\ &= \mathbf{d}_A^T (\boldsymbol{\beta} \boldsymbol{\beta}^T) \mathbf{d}_A - \mathbf{d}_B^T (\boldsymbol{\beta} \boldsymbol{\beta}^T) \mathbf{d}_B \\ &\quad + \sum_{i \in S} (w_{Ai} - 1)^2 \sigma_i^2 - \sum_{i \in S} (w_{Bi} - 1)^2 \sigma_i^2. \end{aligned} \quad (7)$$

To estimate Δ , we first consider how to estimate $\boldsymbol{\beta}$ and the variance of $\hat{\boldsymbol{\beta}}$. The usual weighted least squares estimator is $\hat{\boldsymbol{\beta}} = S_x^{-1} S_{xy}$ where $S_x = \sum_{i \in S} v_i^{-1} \mathbf{x}_i \mathbf{x}_i^T$ and $S_{xy} = \sum_{i \in S} v_i^{-1} \mathbf{x}_i Y_i$. The usual estimator of the variance of $\hat{\boldsymbol{\beta}}$ is $\widehat{\text{var}}[\hat{\boldsymbol{\beta}}] = \hat{\sigma}^2 S_x^{-1}$ where $\hat{\sigma}^2 = \sum_{i \in S} (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i)^2 v_i^{-1} / (n - p)$.

We can estimate $(\boldsymbol{\beta} \boldsymbol{\beta}^T)$ unbiasedly by $(\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}[\hat{\boldsymbol{\beta}}])$. Hence the following is an unbiased estimator of Δ :

$$\begin{aligned} \hat{\Delta} &= \mathbf{d}_A^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}[\hat{\boldsymbol{\beta}}]) \mathbf{d}_A - \mathbf{d}_B^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}[\hat{\boldsymbol{\beta}}]) \mathbf{d}_B \\ &\quad + \sum_{i \in S} (w_{Ai} - 1)^2 \hat{\sigma}^2 v_i - \sum_{i \in S} (w_{Bi} - 1)^2 \hat{\sigma}^2 v_i. \end{aligned} \quad (8)$$

Expression (7) applies, and estimator (8) is an unbiased estimator of it, for any weighted predictors \hat{T}_A and \hat{T}_B . We are concerned with the special case where \hat{T}_A and \hat{T}_B are BLUPs. In this case, \hat{T}_A is calibrated to \mathbf{T}_x so that $\mathbf{d}_A = \mathbf{0}$, and so (8) simplifies to

$$\begin{aligned} \hat{\Delta} &= -\mathbf{d}_B^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}[\hat{\boldsymbol{\beta}}]) \mathbf{d}_B \\ &\quad + \sum_{i \in S} (w_{Ai} - 1)^2 \hat{\sigma}^2 v_i - \sum_{i \in S} (w_{Bi} - 1)^2 \hat{\sigma}^2 v_i. \end{aligned} \quad (9)$$

3.3 An important special case

In this Subsection, we make the assumptions stated in Section 3.2, and further assume that $v_i = 1$ for all i . We also assume that the dimension of \mathbf{x}_{2i} is 1, *i.e.* that we are considering whether or not one particular auxiliary variable from \mathbf{x}_i is to be used in prediction. Expressions (7) and (9) simplify in this case.

Let u_i be the residual of a regression of x_{2i} on \mathbf{x}_{1i} :

$$u_i = x_{2i} - C^T \mathbf{x}_{1i}$$

$$C = \left(\sum_{i \in S} \mathbf{x}_{1i} \mathbf{x}_{1i}^T \right)^{-1} \sum_{i \in S} \mathbf{x}_{1i} x_{2i}.$$

Using straightforward linear algebra operations, it can be shown that

$$\boldsymbol{\beta}^T \mathbf{d}_B = -\beta_2 \sum_{i \in r} u_i$$

and that

$$\sum_{i \in S} (w_{Ai} - 1)^2 - \sum_{i \in S} (w_{Bi} - 1)^2 = \left(\sum_{i \in r} u_i \right)^2 S_u^{-1}$$

where $S_u = \sum_{i \in S} u_i^2$.

Hence (7) becomes

$$\Delta = \sigma^2 \left(\sum_{i \in r} u_i \right)^2 S_u^{-1} - \left(\sum_{i \in r} u_i \right)^2 \beta_2^2.$$

CSW show that $d_B^T \widehat{\text{var}}[\hat{\boldsymbol{\beta}}_2] d_B = \hat{\sigma}^2 (\sum_{i \in r} u_i)^2 S_u^{-1}$. Hence (9) becomes

$$\hat{\Delta} = \left(\sum_{i \in r} u_i \right)^2 (2\hat{\sigma}^2 S_u^{-1} - \hat{\beta}_2^2).$$

It is proposed that \hat{T}_A be adopted when $\hat{\Delta} < 0$, and \hat{T}_B be used otherwise. It follows that we adopt \hat{T}_A whenever $\hat{\beta}_2^2 > 2\hat{\sigma}^2 S_u^{-1}$. As noted by CSW, this is equivalent to adopting \hat{T}_A whenever $F = \hat{\beta}_2^2 / (\hat{\sigma}^2 S_u^{-1})$ is greater than 2. Notice that F is the usual F-statistic for testing the null hypothesis that $\beta_2 = 0$. For large n , the cutoff for the F-test at the 5% significance level is 3.96, whereas we have arrived at a cutoff of 2 for adopting the larger set of variables. Thus, the decision to use A instead of B on the basis of a test of significance requires more evidence against B than a simple comparison of the estimated MSEPs of \hat{T}_A and \hat{T}_B would suggest. That is, using $\hat{\Delta}$ leads to larger models compared to using significance testing.

3.4 Heteroscedasticity-robust estimation of Δ

The estimators of Δ in Sections 3.2 and 3.3 relied on knowing $\text{var}[Y_i]$ at least up to a constant of proportionality. In practice, variances are at best known approximately, and methods which do not rely on an assumption of known variance may perform better. We will use an estimator of σ_i^2 which, assuming model (5), is approximately unbiased for σ_i^2 in general, and exactly unbiased if $\sigma_i^2 = \sigma^2$:

$$\hat{\sigma}_i^2 = \frac{n}{n-p} (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i)^2.$$

(An alternative estimator would be $\hat{\sigma}_i^2 = (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i)^2$, as in Royall and Cumberland 1981b.)

The estimator of $\boldsymbol{\beta}$ would still be the weighted least squares estimator given by (3). The variance of $\hat{\boldsymbol{\beta}}$ is

$$\begin{aligned} \text{var}[\hat{\boldsymbol{\beta}}] &= \text{var}[S_x^{-1} S_{xy}] \\ &= \text{var} \left[S_x^{-1} \sum_{i \in S} \mathbf{x}_i Y_i \right] \\ &= S_x^{-1} \left(\sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^T \sigma_i^2 \right) S_x^{-1}. \end{aligned}$$

This can be estimated by

$$\widehat{\text{var}}_{\text{robust}}[\hat{\boldsymbol{\beta}}] = S_x^{-1} \left(\sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^T \hat{\sigma}_i^2 \right) S_x^{-1}.$$

Hence we can estimate Δ by

$$\begin{aligned} \hat{\Delta} &= \mathbf{d}_A^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}_{\text{robust}}[\hat{\boldsymbol{\beta}}]) \mathbf{d}_A \\ &\quad - \mathbf{d}_B^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}_{\text{robust}}[\hat{\boldsymbol{\beta}}]) \mathbf{d}_B \\ &\quad + \sum_{i \in S} (w_{Ai} - 1)^2 \hat{\sigma}_i^2 - \sum_{i \in S} (w_{Bi} - 1)^2 \hat{\sigma}_i^2. \end{aligned} \quad (10)$$

3.5 Estimation of Δ in multi-stage sampling

The estimators of Δ in Sections 3.2, 3.3 and 3.4 all assumed that the values of Y are independent for different units. In multi-stage sampling, a sample of primary sampling units (PSUs) is initially selected. A sample of units within the selected PSUs is then selected. For example, PSUs may be areas and units may be households or people; or PSUs could be schools and units could be students. Typically units from the same PSUs tend to be similar, so that values of Y_i and Y_j may be correlated if i and j belong to the same PSU. This section develops an estimator of Δ which is approximately unbiased even when there are correlations between values of Y within the same PSU.

Let s_I be the sample of PSUs, selected from the population U_I . Let s_g be the sample of units from PSU g , where $g \in s_I$. Let $r_I = U_I - s_I$ and $r_g = U_g - s_g$. We assume model (5), and further assume that Y_i and Y_j are uncorrelated for $i \in g_1$ and $j \in g_2$ if $g_1 \neq g_2$. The values Y_i and Y_j may be correlated if $i \neq j$ with $i, j \in U_g$.

Let $\hat{T} = \sum_{i \in S} w_i Y_i$ be any predictor and let $\mathbf{d} = \sum_{i \in S} w_i \mathbf{x}_i - \mathbf{T}_x$. The bias of \hat{T} is $\boldsymbol{\beta}^T \mathbf{d}$, as in Section 3.2. The variance of $(\hat{T} - T_Y)$ is

$$\begin{aligned} \text{var}[\hat{T} - T_Y] &= \text{var} \left[\sum_{i \in S} (w_i - 1) Y_i - \sum_{i \in r} Y_i \right] \\ &= \text{var} \left[\sum_{g \in s_I} \left(\sum_{i \in s_g} (w_i - 1) Y_i - \sum_{i \in r_g} Y_i \right) - \sum_{g \in r_I} \sum_{i \in U_g} Y_i \right] \\ &= \sum_{g \in s_I} \text{var} \left[\sum_{i \in s_g} (w_i - 1) Y_i - \sum_{i \in r_g} Y_i \right] + \sum_{g \in r_I} \text{var} \left[\sum_{i \in U_g} Y_i \right]. \end{aligned}$$

It is further assumed that the variance of $\sum_{i \in r_g} Y_i$ and the covariance between $\sum_{i \in r_g} Y_i$ and $\sum_{i \in s_g} (w_i - 1) Y_i$ are negligible relative to other terms. This is the case if cluster sampling is used (because in this case $s_g = U_g$ and r_g is empty) or if the sampling fraction within PSUs is small. The variance becomes

$$\text{var} [\hat{T} - T_Y] \approx \sum_{g \in s_I} \text{var} \left[\sum_{i \in s_g} (w_i - 1) Y_i \right] + \sum_{g \in r_I} \text{var} \left[\sum_{i \in U_g} Y_i \right].$$

Applying this to Δ , we get:

$$\begin{aligned} \Delta &= \text{MSEP}[\hat{T}_A] - \text{MSEP}[\hat{T}_B] \\ &= \mathbf{d}_A^T (\boldsymbol{\beta} \boldsymbol{\beta}^T) \mathbf{d}_A - \mathbf{d}_B^T (\boldsymbol{\beta} \boldsymbol{\beta}^T) \mathbf{d}_B + \sum_{g \in s_I} \text{var} \left[\sum_{i \in s_g} (w_{Ai} - 1) Y_i \right] \\ &\quad - \sum_{g \in s_I} \text{var} \left[\sum_{i \in s_g} (w_{Bi} - 1) Y_i \right]. \end{aligned} \quad (11)$$

To estimate Δ , we need estimators of the variance of $\hat{\boldsymbol{\beta}}$, and of $(\sum_{i \in s_g} (w_i - 1) Y_i)$.

Firstly, notice that

$$\begin{aligned} \text{var} [\hat{\boldsymbol{\beta}}] &= \text{var} \left[S_x^{-1} \sum_{g \in s_I} \sum_{i \in s_g} \mathbf{x}_i Y_i \right] \\ &= S_x^{-1} \sum_{g \in s_I} \text{var} \left[\sum_{i \in s_g} \mathbf{x}_i Y_i \right] S_x^{-1}. \end{aligned}$$

This can be estimated using the “ultimate cluster variance” method by

$$\widehat{\text{var}}_{\text{ucv}} [\hat{\boldsymbol{\beta}}] = S_x^{-1} \sum_{g \in s_I} \left(\sum_{i \in s_g} \mathbf{x}_i (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i) \right)^2 S_x^{-1}.$$

This is a well known estimator of the variance of a weighted sum from clustered data, and is equivalent to Valliant *et al.* (2000, 9.5.5, page 312). The variance has been called a “sandwich-level variance estimator using the cluster-level residuals” (Valliant *et al.* 2000) and an “ultimate cluster variance” (*e.g.*, Wolter 1985 describes essentially the same idea in a randomization framework).

The variance of $(\sum_{i \in s_g} (w_i - 1) Y_i)$ can also be estimated by the ultimate cluster variance method:

$$\widehat{\text{var}} \left[\sum_{i \in s_g} (w_i - 1) Y_i \right] = \left\{ \sum_{i \in s_g} (w_i - 1) (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i) \right\}^2.$$

Hence we can estimate Δ by

$$\begin{aligned} \hat{\Delta}_{\text{ucv}} &= \mathbf{d}_A^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}_{\text{ucv}} [\hat{\boldsymbol{\beta}}]) \mathbf{d}_A \\ &\quad - \mathbf{d}_B^T (\hat{\boldsymbol{\beta}} \hat{\boldsymbol{\beta}}^T - \widehat{\text{var}}_{\text{ucv}} [\hat{\boldsymbol{\beta}}]) \mathbf{d}_B \\ &\quad + \sum_{g \in s_I} \left\{ \sum_{i \in s_g} (w_{Ai} - 1) (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i) \right\}^2 \\ &\quad - \sum_{g \in s_I} \left\{ \sum_{i \in s_g} (w_{Bi} - 1) (Y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i) \right\}^2. \end{aligned} \quad (12)$$

4. Simulation study

4.1 Simulation of farm survey

Population and sampling scheme

The population distribution of the auxiliary variables, the sample and population size, and heteroscedasticity and other properties of the variable of interest would all be expected to play a part in the performance of the adaptive BLUPs. To make a realistic assessment of the performance of these estimators, a simulation study based on a large, realistic population is needed.

We generated a simulation population of 80,000 units, using sample data on 1652 farms from the 1988 Australian Agricultural and Grazing Industry Survey (AAGIS) as a starting point. Total cash crop was used as the survey variable of interest, and potential auxiliary variables included DSE (a derived size estimate), number of sheep, crops area, number of beef cattle, region (29 regions) and industry (5 industries). DSE was a linear combination of the sheep, crops area and beef cattle variables. The dataset also contained a sampling weight which was approximately equal to the inverse of the selection probability. 27 outliers with very large values of DSE were removed, as these would normally be placed in a completely enumerated stratum in a survey. A population of 80,000 was then constructed by probability proportional to size sampling with replacement, with probabilities proportional to the estimation weight on the original sample file.

250 samples were then selected without replacement from the simulation population. The samples were stratified by Region and DSE, with DSE divided into four categories, to give 116 strata. The category boundaries were set such that the category sums of DSE were equal. Total sample sizes of 250, 500, 1,000 and 1,500 were simulated. The stratum sample sizes were proportional to the original AAGIS sample sizes by Region and DSE.

Auxiliary variables and stepwise selection method

Auxiliary variables were included corresponding to the model containing: an intercept; sheep (x1); crops area (x2); beef cattle (x3); Industry; interaction of Industry and x1, x2 and x3; and Region. This gives a total of 52 potential auxiliary variables. Some of these variables are collinear, but are still included in the set of potential variables, to give the model selection process a wider choice of possible models. We also considered the set of 139 auxiliary variables which included this set as well as the interaction of Region and x1, x2 and x3. Models were constructed by forward selection starting with the intercept-only model. Variables were added based on which step most reduced the estimated MSEP, for several alternative estimators of Δ .

Stepwise selection was also trialled but was substantially slower to run for the larger variable set, and did not greatly improve the efficiency of the adaptive BLUPs.

An adaptive BLUP was also calculated based on statistical significance, with $p < 0.05$ being the cutoff for inclusion. For each progressive model, the statistical significance of adding each of the variables not in the model was assessed, using a standard t-test. The variable with the lowest p-value was included in the model at each step. When there were no further significant variables which could be added, the procedure terminated and this was the model chosen.

A number of modifications were needed for the forward selection algorithm to work reliably: auxiliary variables were not added to the model if they had a pairwise Pearson correlation of 0.95 or higher (or -0.95 or lower); and variables were not added if this would result in the calibration equations not being solvable.

Estimators used

Several BLUPs were calculated: with all auxiliary variables included; with just Intercept and DSE; and with auxiliary variables chosen by forward selection using the non-robust estimator of Δ (described in Section 3.2) or the heteroscedasticity-robust estimator of Δ (described in Section 3.4), from either the set of 52 or the set of 139 potential auxiliary variables. (The larger set of 139 variables was only evaluated for sample sizes of 500 and above.)

Ridge estimators (e.g., Bardsley and Chambers 1984) are an alternative approach to the problem of variable selection, so we included them in the simulation to compare their performance to that of the adaptive BLUPs. The estimators we have so far considered either include or exclude each variable. If a variable is included, then the weights must calibrate on that variable exactly, in the sense that $\sum_{i \in s} w_i x_i = T_x$. Ridge regression introduces a penalty for non-calibration, but does not necessarily require that the weights provide perfect calibration for all variables. In ridge regression, the vector of sample weights w is chosen to minimise

$$\sum_{i \in s} (w_i - 1)^2 v_i^{-1} + \sum_{j=1}^p c_j^{-1} \left(\sum_{i \in s} w_i x_{ij} - T_{xj} \right)^2.$$

The c_j are non-negative cost coefficients indicating the priority to be placed on meeting calibration constraint j . A value of 0 indicates that the constraint must be met precisely and larger cost coefficients result in placing less weight on the constraint. Thus the ridge estimator allows for a smooth reduction in the effective dimension of the model, by effectively interpolating between including a calibration variable ($c_j = 0$) and excluding it ($c_j = \infty$).

Typically the c_j are set to λc_j^* , where c_j^* reflect a somewhat subjective assessment of the relative importance of each constraint, and λ is chosen to ensure that the final weights w_i have reasonable properties, for example are all greater than or equal to 0, or to 1. We set c_j^* to 0 for the constant (reflecting an intercept in the model), to 1 for x_1, x_2 and x_3 , to 10 for the region indicators, to 5 for the industry indicators, and to 100 for interactions. The choice of c_j^* was based on which variables were thought to be likely to be most useful. The value of λ was numerically determined for each sample to be the smallest value such that all weights were greater than or equal to 1.

All of the methods were based on the same procedure for modelling $\text{var}_M[Y_i]$. Firstly, a simple model with the intercept, x_1, x_2 and x_3 was fitted to the sample values of Y using ordinary least squares. The log of the squared residuals from this model were then regressed against the log of DSE. The fitted values of this model were raised to the power of e to give estimates of σ_i^2 for each $i \in s$. The estimated values of σ_i^2 were then truncated so that no values were more than 4 times, or less than one quarter, of the median value. This adjustment was made to avoid extreme values of σ_i^{-2} which might lead to instability in calculating weighted least squares estimates of β . Results were somewhat sensitive to the variance modelling procedure, particularly the final adjustment to avoid extreme values: BLUPs based on a crude variance model with $\sigma_i^2 \propto \text{DSE}_i$ had variances around 10-20% higher than the BLUPs shown here.

Results

Table 1 shows the Relative Root Mean Squared Error (RRMSE) of the various calibrated predictors. The first four rows of the table are for the first set of auxiliary variables (52 potential variables) and the last three rows are for the second set (139 potential auxiliary variables). Biases are not shown but were generally a relatively small component of the mean squared error for all of the predictors shown, except for the BLUP based on an intercept and DSE model, which was quite biased. This was somewhat surprising as we expected that a good trade-off between bias and variance would imply that biases were a non-negligible component of the mean squared error. Details on the biases and relative variances of the predictors can be found in Tables A1 and A2 of Clark and Chambers (2008).

Of the adaptive BLUPs, the significance criteria performed the best in all cases, followed by the nonrobust criteria, with the robust criteria performing worst. For the smaller set of 52 potential variables, the adaptive BLUPs based on the nonrobust and significance criteria performed better than the nonadaptive BLUPs for $n = 250$ and $n = 500$; for $n = 1,000$ and $n = 1,500$, they performed

slightly worse than the BLUP with all variables but better than the intercept and size BLUP. For the larger set of 139 potential variables, the adaptive BLUPs based on the nonrobust and significance criteria performed better than the nonadaptive BLUPs for all sample sizes, particularly for smaller values of n .

Table 1
RRMSE (%) of AAGIS predictors of total cash crops

# Vars	n	BLUP		Adaptive BLUP		Ridge	
		all	int+size	nonrobust $\hat{\Delta}$	robust $\hat{\Delta}$	Sig.Test	
52	250	3.59	3.02	2.97	3.09	2.87	3.30
	500	2.35	2.54	2.33	2.33	2.30	2.31
	1,000	1.56	2.21	1.58	1.64	1.57	1.54
	1,500	1.36	2.22	1.39	1.41	1.37	1.37
139	500	3.52	2.54	2.99	3.44	2.29	2.27
	1,000	1.77	2.21	1.75	1.92	1.72	1.59
	1,500	1.56	2.22	1.51	1.64	1.42	1.42

The Ridge estimator generally performed about as well as the best of the adaptive BLUPs when there were 52 auxiliary variables, and slightly better when there were 139 potential variables.

Table 2 shows how many auxiliary variables were selected for the two adaptive BLUPs. The robust $\hat{\Delta}$ led to larger sets of auxiliary variables than the non-robust, with about 10 more auxiliary variables selected. The significance criteria led to even smaller variable sets (6-10 less variables than from the non-robust criteria).

Table 2
Mean (Interquartile range) of number of auxiliary variables selected in AAGIS

# Vars	n	nonrobust $\hat{\Delta}$	robust $\hat{\Delta}$	Sig.Test
52	250	16.0 (14.0-18.0)	26.9 (24.0-29.0)	9.6 (8.0-11.0)
	500	18.6 (16.0-21.0)	27.4 (25.0-30.0)	11.5 (10.0-13.0)
	1,000	23.6 (21.0-26.0)	29.6 (26.0-33.0)	14.4 (13.0-16.0)
	1,500	27.3 (25.0-29.0)	32.3 (30.0-35.0)	17.2 (16.0-18.8)
139	500	42.1 (37.0-47.0)	69.4 (62.0-75.0)	23.2 (21.0-26.0)
	1,000	51.5 (47.0-56.0)	74.2 (69.0-79.8)	29.9 (27.0-33.0)
	1,500	59.2 (55.0-64.0)	75.8 (71.0-81.0)	34.9 (32.0-38.0)

Table 3 shows the confidence interval (CI) non-coverage of the various predictors. 90% CIs were defined as the estimator ± 1.64 standard errors, where the variance was estimated using a heteroscedasticity-robust variance estimator (Royall and Cumberland 1978). Following common practice, CIs were based on estimated variance not estimated mean squared error of prediction. The simulation estimates of the non-coverage rates are fairly rough given that only 250 simulations were used. A larger simulation study could be used to give more precise estimates of coverage, but this was not pursued due to the

computationally intensive nature of the stepwise selection process. Table 3 suggests that: the BLUP using just intercept plus size had high non-coverage as did the adaptive BLUP based on robust $\hat{\Delta}$. The other estimators generally had non-coverage rates close to the nominal 10%.

Table 3
Confidence interval non-coverage in AAGIS

# Vars	n	BLUP		Adaptive BLUP			Ridge
		all	int+size	nonrobust $\hat{\Delta}$	robust $\hat{\Delta}$	Sig.Test	
52	250	10.0	6.4	10.4	16.8	11.2	10.0
	500	8.0	13.2	12.0	17.2	10.8	8.0
	1,000	7.6	20.4	9.2	12.0	8.4	8.4
	1,500	8.8	34.8	9.2	13.2	9.6	8.8
139	500	16.8	13.2	18.0	29.2	12.8	8.8
	1,000	12.4	20.4	14.0	20.4	13.2	7.2
	1,500	13.6	34.8	13.6	19.6	12.4	11.2

Total cash crops is a major variable of interest in the AAGIS survey, but the totals of other variables are also important, including Farm Equity. For practical reasons, a single set of weights is normally used for all variables. Table 4 shows how well the adaptive calibration weights designed for the Total Cash Crops (TCC) variable performed when used to estimate the total of Farm Equity. For the case of 52 potential auxiliary variables, the adaptive BLUP weights chosen based on TCC (using non-robust $\hat{\Delta}$) performed reasonably well, as did the ridge estimator. Improvements could be made, however, by choosing auxiliary variables based on Equity.

Table 4
RRMSE (%) of AAGIS predictors of total equity

# Vars	n	BLUP		Adaptive BLUP (nonrobust $\hat{\Delta}$)		Ridge
		all	int+size	based on TCC	based on Equity	
52	250	6.85	6.45	6.51	6.13	6.78
	500	4.44	4.44	4.61	4.40	4.28
	1,000	3.09	3.12	3.42	3.14	3.10
	1,500	2.54	2.58	2.90	2.58	2.54
139	500	5.53	4.93	4.98	4.74	4.20
	1,000	3.68	4.03	3.23	3.15	3.08
	1,500	3.04	3.63	2.66	2.60	2.57

4.2 Simulation of Labour Force Survey

Population and Sampling Scheme

A simulation population was constructed by selecting a simple random sample without replacement of 30,000 people aged 15-64 from the 1% sample file of the 1991 Australian Census of Population and Housing. The variable of interest was Employment (1 for employed people, 0 for others). The simulation population was divided into

simulated primary sampling units (PSUs) containing 75 people each, in such a way that the intra-cluster correlation was 0.05. (This is a fairly typical intra-class correlation for the employment variable within primary sampling units in a household survey. See for example Clark and Steel 2002). The algorithm for defining clusters was to sort the data by a randomly generated $N(0, \gamma^2)$ variable plus the employment variable, then to define clusters as sequential sets of 75 people, where γ was chosen so as to give the desired intra-cluster correlation.

The simulation consisted of 250 repeated two-stage samples. The first stage was a simple random sample without replacement of m PSUs and the second stage was a simple random sample without replacement of 20 people from each selected PSU. The total sample size was set to be $n = 200, 400$ and 1,000 people. Most national household surveys have sample sizes much larger than this, but it is common to construct estimation post-strata within states or provinces, and the sample sizes for these areas would often be in the range 200-1,000.

The potential auxiliary variables were age by sex, where age was recorded in single years for 16-24 year olds, then in five year age groups 25-29, 30-34, ..., 55-59 year olds, and 60+ year olds.

Non-response

One of the main reasons why age and sex are used as auxiliary variables in household surveys is that non-response is known to depend on age and sex. For example, young men are typically the group with the lowest response rates. Non-response was simulated by assuming that the logit of the probability of response was equal to $1.8 - ((age - 50)/25)^2$ for men, and $2 - 0.7((age - 50)/25)^2$ for women. This model gave a response rate of 75%. The initial sample size was increased so that the final responding sample size was equal to $n = 200, 400$ or 1,000.

Auxiliary variables and stepwise selection method

The potential auxiliary variables were based on age by sex cells. The definition of the x -variables is shown in Table 5. This parameterization was chosen so that the auxiliary variables corresponding to specific ages or agegroups can be dropped while still giving a sensible model. For example, if all auxiliary variables were included except for x_{4i} , then the model expected value for people aged 17 would be the same as those aged 16, rather than being equal to the intercept parameter. Even better results might be obtained from using more sophisticated parameterizations such as spline models and this will be investigated in a future study.

Table 5
Potential auxiliary variables in labour force survey simulation

Variable	Definition
x_{1i}	1 (corresponding to intercept in model for Y)
x_{2i}	1 if person i male -1 if female
x_{3i}	1 if person i aged 16 or over
x_{4i}	1 if person i aged 17 or over
\vdots	\vdots
$x_{12,i}$	1 if person i aged 25 or over
$x_{13,i}$	1 if person i aged 30 or over
\vdots	\vdots
$x_{19,i}$	1 if person i aged 60 or over
$x_{20,i}$	x_{3i} if person i male $-x_{3i}$ if female
\vdots	\vdots
$x_{36,i}$	$x_{19,i}$ if person i male $-x_{19,i}$ if female

Stepwise selection was used to select variables, starting with the intercept-only model. At each step, variables could be added or removed, according to which gave the best reduction in the criteria. If the stepwise selection began cycling (for example, adding x_1 , then adding x_2 , then removing x_1 , then removing x_2 , then adding x_1 , etc), then the model building process stopped, and the the current model was used as the final model. The estimators of Δ used were the non-robust estimator, the robust (to heteroscedasticity) estimator and the ultimate cluster variance (UCV) estimator which is robust to heteroscedasticity and correlations within PSUs. Significance tests were not used as they would need to incorporate correlations within PSUs to be realistic. Results for the ridge estimator are not shown because negative weights rarely occurred in this simulation, so that this estimator performed very similarly to the BLUP using all auxiliary variables.

Results

Table 6 shows the RRMSE of the various adaptive and non-adaptive BLUPs. There was relatively little difference in RRMSE between the BLUP with intercept only and the BLUP with all auxiliary variables. It is therefore not surprising that at best minor gains were made by using the adaptive BLUPs rather than using the BLUP with all variables. The adaptive BLUP using the non-robust $\hat{\Delta}$ gave the lowest RRMSE in all cases.

Table 6
RRMSE of labour force survey predictors of employment

n	BLUP		Adaptive BLUP		
	all	intercept	nonrobust $\hat{\Delta}$	robust $\hat{\Delta}$	UCV $\hat{\Delta}$
200	6.54	6.77	6.44	7.06	6.96
400	4.72	4.76	4.61	4.72	4.65
1,000	2.45	2.70	2.43	2.45	2.49

Table 7 shows the mean number of variables selected for each of the adaptive BLUPs. Of the 36 potential auxiliary variables, between about 5 and 9 variables were selected based on the non-robust $\hat{\Delta}$. The number of variables selected increased as the sample size increased. The heteroscedasticity-robust criterion resulted in larger sets of auxiliary variables, and the UCV criterion gave even larger sets.

Table 7
Mean (Interquartile range) of number of auxiliary variables selected in labour force simulation

n	Variable Selection Method		
	nonrobust	robust	UCV
200	6.5 (5.0- 8.0)	13.4 (10.0-16.0)	16.1 (13.0-19.0)
400	7.4 (6.0- 8.0)	12.1 (9.0-15.0)	14.5 (12.0-17.0)
1,000	8.6 (7.0-10.0)	11.6 (10.0-13.0)	14.2 (12.0-17.0)

Table 8 shows the confidence interval (CI) non-coverage of the various predictors. 90% CIs were defined as the estimator ± 1.64 standard errors, where the variance was estimated using a UCV variance estimator. Table 8 shows that the BLUP using all auxiliary variables had high non-coverage for $n = 200$ and 400. The adaptive BLUP using nonrobust $\hat{\Delta}$ had reasonably close to nominal coverage, while the other adaptive BLUPs had high non-coverage.

Table 8
Confidence interval non-coverage (%) for predictors of employment

n	BLUP		Adaptive BLUP		UCV $\hat{\Delta}$
	all	intercept	nonrobust $\hat{\Delta}$	robust $\hat{\Delta}$	
200	17.6	12.0	12.0	20.0	24.0
400	17.2	12.0	14.8	16.8	17.6
1,000	6.4	11.6	7.6	6.8	9.6

Table 9 shows how well the various weights performed when used to estimate a different variable, unemployment (equal to 1 for unemployed people and 0 otherwise). Adaptive BLUPs were calculated using the non-robust $\hat{\Delta}$, with the variable of interest given by Employment, and by Unemployment. The adaptive BLUP with variables chosen for Employment had RRMSE between the non-adaptive BLUP with all variables and the non-adaptive BLUP with intercept only. This suggests that this adaptive BLUP gives reasonable results even when applied to variables other than employment. The adaptive BLUP based on Unemployment actually had higher RRMSE. This may be because the auxiliary variables had little or no predictive power for unemployment, so that attempting to tailor the choice of auxiliary variables for this variable of interest did not work well.

Table 9
RRMSE of labour force survey predictors of unemployment

n	BLUP		Adaptive BLUP	
	all	intercept	based on emp	based on unemp
200	36.3	32.6	34.5	36.0
400	24.1	21.7	22.8	23.7
1,000	14.5	14.2	14.1	14.2

5. Discussion

The simulation studies described here showed that adaptive BLUPs can give useful gains compared to simple non-adaptive alternatives. In both the farm survey and the labour force survey simulations, the adaptive BLUPs based on a nonrobust estimator of Δ and based on significance testing both had lower MSEP than non-adaptive estimators in almost all cases. In the case of the farm survey, the gains were sometimes substantial compared to either always using the full model or always using the intercept plus size variable model. In the case of the labour force survey, the gains were minor. The adaptive BLUPs also gave reasonable confidence interval coverage.

The adaptive BLUPs based on the robust and UCV criteria performed much worse than the other adaptive BLUPs. This is surprising, as the AAGIS data is known to be heteroscedastic and the Labour Force data was clustered suggesting that the UCV criteria should have given good results. Further analysis of the farms survey simulation showed that $\hat{\Delta}_{\text{robust}}$ had higher variances than $\hat{\Delta}_{\text{nonrobust}}$ in the great majority of cases, particularly for auxiliary variables with little predictive power - see the Appendix of Clark and Chambers 2008 for details. This suggests that the robust method would tend to select counter-productive auxiliary variables more often and could explain its poor performance.

There was a general tendency for all of the adaptive procedures to choose too many auxiliary variables, but despite this, the adaptive estimators generally performed better than or similar to simple non-adaptive alternatives. We suggest that in practice, an automatic model search (using either a non-robust $\hat{\Delta}$ or a statistical significance criterion) should be used in conjunction with some subjective judgement. For example, models could be selected from several sets of potential auxiliary variables of different sizes. If the larger sets gave only small apparent improvements, then the statistician might decide to restrict to a smaller set, even if apparently slightly suboptimal.

Ridge estimators also performed reasonably well in terms of RRMSE and confidence interval coverage. They generally gave similar results to the adaptive BLUPs for estimating the total of the variable of interest when the choice of auxiliary variables was based on this variable.

However, when the adaptive BLUP weights were applied to different variables, the ridge estimators performed slightly better. An even better approach may be adaptively choose both which auxiliary variables to include and how to apply ridging, based on some criterion calculated from the sample. This will be the topic of future research.

One concern that has been raised with the prediction approach to finite population sampling is its non-robustness to the omission of important auxiliary variables. In our simulations from farm economic data and social data, the adaptive predictors had low bias and lower mean-squared error than the non-adaptive estimators in most of the wide range of cases in our simulation study, and were never substantially worse. Provided that all design variables are considered as potential auxiliary variables, adaptive calibration provides a robust and efficient strategy for finite population prediction.

Acknowledgements

This work was jointly supported by the Australian Research Council and the Australian Bureau of Statistics. An associate editor and two referees made detailed and thoughtful comments which considerably improved the paper.

References

- Bardsley, P., and Chambers, R. (1984). Multipurpose estimation from unbalanced samples. *Applied Statistics*, 33(3), 290-299.
- Brewer, K.R.W. (1963). Ratio estimation and finite populations: Some results deductible from the assumption of an underlying stochastic process. *Australian Journal of Statistics*, 5, 93-105.
- Chambers, R., Skinner, C. and Wang, S. (1999). Intelligent calibration. *Bulletin of the International Statistical Institute*, 58(2), 321-324.
- Clark, R.G., and Chambers, R.L. (2008). *Adaptive calibration for prediction of finite population totals*. University of Wollongong Centre for Statistical and Survey Methodology. Available from <http://www.cssm.uow.edu.au>.
- Clark, R.G., and Steel, D.G. (2002). The effect of using household as a sampling unit. *International Statistical Review*, 70(2), 289-314.
- Deville, J.-C., and Särndal, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association*, 87(418), 376-382.
- Hansen, M., Madow, W. and Tepping, B. (1983). An evaluation of model-dependent and probability-sampling inferences in sample surveys. *Journal of the American Statistical Association*, 78, 776-793.
- Royall, R.M. (1970). On finite population sampling theory under certain linear regression models. *Biometrika*, 57(2), 377-387.
- Royall, R.M. (1976). The linear least squares prediction approach to two-stage sampling. *Journal of the American Statistical Association*, 71(355), 657-664.
- Royall, R.M., and Cumberland, W.G. (1978). Variance estimation in finite population sampling. *Journal of the American Statistical Association*, 73(362), 351-358.
- Royall, R.M., and Cumberland, W.G. (1981a). An empirical study of the ratio estimator and estimators of its variance. *Journal of the American Statistical Association*, 76(373), 66-80.
- Royall, R.M., and Cumberland, W.G. (1981b). The finite-population linear regression estimator and estimators of its variance - an empirical study. *Journal of the American Statistical Association*, 76(376), 924-930.
- Royall, R.M., and Herson, J. (1973a). Robust estimation in finite populations 1. *Journal of the American Statistical Association*, 68(344), 880-889.
- Royall, R.M., and Herson, J. (1973b). Robust estimation in finite populations 2: Stratification on a size variable. *Journal of the American Statistical Association*, 68(344), 890-893.
- Särndal, C.-E., Swensson, B. and Wretman, J. (1992). *Model assisted survey sampling*. New York: Springer-Verlag.
- Silva, P.L.N., and Skinner, C. (1997). Variable selection for regression estimation in finite populations. *Survey Methodology*, 23, 23-32.
- Skinner, C., and Silva, P.L.N. (1997). Variable selection for regression estimation in the presence of nonresponse. *Proceedings of the Section on Survey Research Methods*, American Statistical Association, 76-81.
- Smith, T.M.F. (1976). The foundations of survey sampling: a review. *Journal of the Royal Statistical Society, Series A*, 139(2), 183-202.
- Valliant, R., Dorfman, A.H. and Royall, R.M. (2000). *Finite Population Sampling and Inference: A Prediction Approach*. New York: John Wiley & Sons, Inc.
- Wolter, K.M. (1985). *Introduction to variance estimation*. New York: Springer-Verlag.