# Consumer Expenditure Estimation Incorporating Generalized Variance Functions in Hierarchical Bayes Models

Paul E. Hinrichs


## Abstract

This paper investigates models for sampling error that would improve the stability of mean annual consumer expenditure estimates. Hierarchical Bayes models presented consider mean estimates from five interview and diary sources for diary only estimates when interview estimates are also available. The investigation breaks down the problem into improving the generalized variance function estimate of variance through three-digit groupings of universal classification codes, estimation of sampling variance using hierarchical models.


## Introduction

Since 1888, nationwide data has been collected on consumer expenditures. The Bureau of Labor Statistics consumer expenditure survey collects information of family expenditures for goods and services as well as income, savings, demographic, and economic characteristics. The survey was given intermittently until 1979 when the need for more timely data created the need for a continuing survey.

Samples are designed to be representative of the U. S. civilian population. The design can be approximated as a stratified multistage sample with two primary sample units per stratum.

If a consumer unit is selected for the interview survey, a sequence of five interviews is administered. Data are collected once per quarter for five consecutive quarters. The first of the five surveys is essentially a bounding survey in which demographic data and an inventory of goods are collected. The next four interviews ask the consumer unit to recall expenditures for an estimated 60 to 70 percent of the total family expenses. The final interview also includes an annual supplement to obtain additional financial background on the consumer unit. In general, the interview is designed to collect information that a consumer would be able to recall over a three-month period. Of those expenses, 65 percent are direct monthly estimates where as 35 percent use quarterly recall.

The diary is used to collect information on expenditures that a household will incur in a given week. This includes food and beverages, personal care products and services that a family would purchase at home or away from home. Consumer units selected for the diary survey record their purchases for two consecutive weeks using the Record of Daily Expenses questionnaire.

A rotating panel of consumer units is used in the surveys. The results of the survey are compiled to publish annual estimates of expenditure for a number of demographic areas and expenditure classes.

There is considerable overlap in expenditure coverage of these two survey instruments. When data are available from both surveys, the survey that is determined to be more reliable is used. The working question is "Which source estimate is the more reliable source?" The data source for an estimate is published within Bureau of Labor Statistics microdata reports.

After the data are collected and errors corrected, the consumer unit estimate is then weighted to control for selection probability, subsampling method, noninterviews in the collection, and consumer unit composition. The estimate of the variance is calculated using the method of balanced repeated replication. The Bureau of Labor Statistics uses 44 replicates to attain its sampling variance estimate.

Work has been done in the past on using a combined data estimator that uses both information sources. Research on combined estimators by Eltinge, Sukasih and Weber was completed in 2000. An estimator based on a vector of five data sources was investigated. The five sources include three interview sources and two diary sources.

Combined data point estimates of annual mean expenditure can then be found using each source at a fine level of expenditure classification. Current studies use the six-digit Universal Classification Code (UCC) level. There are 665 distinct UCCs used in the diary and interview surveys from 1987 to 1997. Of those, 180 are diary-only UCCs, 224 are interview-only UCCs and 261 are interview and diary UCCs.

For the 261 UCCs that can be estimated from either the interview or the diary, the Bureau of Labor Statistics publishes estimates based entirely on the diary data for 78 of the UCCs. The remaining 183 estimates are based on the interview data. The estimation source is an area of ongoing study within the Bureau of Labor Statistics and can change from year to year. Many factors are weighed in determining which estimate to use in estimating the expenditure. In general, preference is given to the source that is thought to be least bias.

Eltinge, Sukasih and Weber (2000) research found a number of efficient estimators for the UCCs but none of the estimators dominated the others uniformly over all UCCs considered. The model used by Eltinge, Sukasih and Weber is similar to a two-stage least squares estimator from econometrics.

In a related study, Eltinge and Sukasih (2001) investigated the effect of different sampling variance estimates across the five data sources in generalized least squares estimators. The authors considered variance estimators that would be more stable than the sampling variance estimated through the balanced repeated replication method. In particular, they modeled the variance through a generalized variance function. Empirical evidence showed that a two-stage estimate of the variance-covariance matrix could be more stable than the balanced repeated replicate estimate of the variance-covariance matrix.

The following work expands on these ideas using hierarchical Bayes modeling for the point estimator and the variance estimator. The model will group the UCCs by their three-digit classification to borrow strength across UCCs and improve both the estimate and the variance estimate for annual expenditures.

## Variance Modeling

For many of the UCCs in the consumer expenditure survey, a household's expenditure in a reporting week will be zero. This in turn leads to a variance estimate for individual UCCs that will be larger than a level at which the estimate is publishable in a given year. Consequently, we seek some support for the variance estimate based on only the diary source. For the 60 UCCs considered in this study we have the supporting data from the three interview sources. The data source thought to be the least bias is used to estimate the mean expenditure. What we have in the interview is an estimate of mean expenditure and sampling variance from a bias source.

Eltinge, Sukasih and Weber (2000) grouped the data for individual UCC, then modeled over the five survey modes and years. Eltinge and Sukasih (2001) demonstrated that there was a relationship between the mean expenditure, $\hat{X}$, and the standard deviation, $\sigma_{\hat{X}}$, of a UCC in the form of $\sigma_{\hat{X}} = \alpha + \beta\hat{X} + \text{error}$. The study focused on a single UCC and found an $R^2$ value of 0.83. The analysis presented here expands the generalized variance function scope to groups of UCCs rather than individual UCCs.

In this section a model for the sampling variance of the mean expenditure is proposed using a generalized variance function. When using a generalized variance model, the variance of a survey estimator is calculated indirectly through a relationship between the variance and the expected value of its survey estimate,

$$\text{Variance} = f(\text{mean}). \qquad (1)$$

A general variance estimator is used when a survey is large and encompasses a large number of subpopulations (Wolter 1985, Eltinge, Sukasih and Weber 2000, Eltinge and Sukasih 2001, Johnson and King 1987, Valliant 1987, Otto and Bell 1995). As a result, a correspondingly large number of subpopulation means and variances are usually desired. The Generalized Variance Function (GVF) model itself describes the mathematical relationship between the variance and the expectation for the various subpopulations.

Fay and Herriot (1977) noted that the coefficient of variation does not depend on the expected value of the sample. The variance would then increase in direct proportion to the expected value. They conclude that the logarithm transformation stabilizes the variance. Because of this we restrict ourselves to linear and logarithmic models relating the variance or relative variance to the mean. The Box-Cox power transformation would be another way to reduce the family of distributions that satisfy (1) and identify an appropriate transformation (see Neter, Kutner, Nachtsheim Wasserman 1996).

To implement generalized variance function techniques, subpopulations of the survey statistics are formed and different models are applied. Subpopulations can be formed by either natural grouping or by data mining. When using data mining, groups, which demonstrate strong relationships among variables are used but may not have any other justification for their grouping. The consumer expenditure survey data leads naturally to a number of different possible natural groupings. For this study we consider data grouped over years, survey collection mode and/or Universal Classification Codes (UCC) three-digit grouping.

### Modeling the Variance

If a sample variance were computed directly from the diary, its usefulness would be limited because of the many zeros in the data set. As such, this variance will be subject to large errors in estimation. Generalized Variance Function theory says it may be possible to improve the stability of the variance by taking advantage of the relationship between variances and means (Wolter 1985, Eltinge et al. 2000, Eltinge & Sukasih 2001, Johnson & King 1987, Valliant 1987). Wolter (1985) listed a number of GVF models in his text. This research will investigate the GVF models given in Table 1 below.

The consumer expenditure description suggests that the data come from five distinct sources of which three are interview and two are diary sources. Let $i = 1, \ldots, 5$ be an index of the five data sources. In addition, data in the consumer expenditure survey for a particular year, $k$, are collected at very fine level of aggregation. UCCs at this level are collected at the 6-digit classification level. The UCCs can also be classified at higher levels of aggregation. Let $j$ represent the index of the grouping for the three-digit grouping for the particular level of aggregation under investigation.

The UCCs when grouped by three-digit classification codes fall into three distinct groupings: household furnishings and equipment, apparel and services, and entertainment. The UCCs used in this analysis are a subset of the full $n$-digit groupings. Complete groups also contain UCCs that are collected using just one data source. Further information about the classifications can be found in the BLS Handbook and at the Consumer Expenditure web site.

Let $\sigma^2$ be the true sampling variance of a whole sample mean estimator $\hat{X}$ and $V^2$ the relative variance of the estimator. We further define the subscripts of $\sigma^2 = \sigma^2_{ijkl}$ as the variance for the $i$th collection mode for the $j$th group in year $k$. If we say $u_j$ represent the number of UCCs in the $j$th UCC classification group then $l = 1, \ldots, u_j$ would be the index of the $l$th UCC in the classification group.

The data consist of whole sample mean annual estimates of expenditure and forty-four balanced repeated replicates of annual expenditure for each UCC and each collection method. A $5 \times 5$ estimate of the variance-covariance matrix was computed using the replicates from the five data sources using the BRR method.

Models

Past research using generalized variance functions yields a rich class of models (Wolter 1985). The models that are investigated here try to capture the major effects when modeling variance or relative variance by the mean.

A scatterplot of estimated relative variances and means indicate that some remediation of the data is necessary. The model error would indicate that the variance was heteroscedastic. In an effort to remediate the model error we apply the logarithm to both mean and variance. Applying a logarithm to this model would further restrict the sample size since all UCCs with a mean estimate of 1 would need to be removed from the model.

The model, $V^2_{ijkl} = \alpha_{ijk} + \beta_{ijk}\hat{X}^{-1}_{ijkl} + \epsilon$, basically says that the relative variance is a quadratic function of the mean estimate with an error term that is weighted by the square of the mean estimate. A form of this quadratic model has been used in a number of Current Population Survey (CPS) estimates of variance. Model $(ii)$ is an implementation of a remedial measure to make error terms normal. Model $(iii)$ is another model implemented in CPS research relating inverse of the relative variance and a linear function of the mean expenditure estimate. The last model is the one model implemented in the research by Eltinge and Sukasih (2001).

Least squares regressions with categorical variables were used to fit the data. The collection method, year and grouping are used as possible categorical variables. With the various possible categorical variables, the coefficients subscripts in the model will vary. For instance, if the coefficients are assumed to be time invariant then the $k$ subscript would be omitted from the model above.

Table 2 gives the results of regression models applied to the consumer expenditure data. The column $ijk$ indicates

Table 1: Models to be applied to Consumer Expenditure Data

| | Model |
|---|---|
| $(i)$ | $V^2_{ijkl} = \alpha_{ijk} + \beta_{ijk}\hat{X}^{-1}_{ijkl} + \epsilon$ |
| $(ii)$ | $\log \sigma^2_{ijkl} = \alpha_{ijk} + \beta_{ijk} \log \hat{X}_{ijkl} + \epsilon$ |
| $(iii)$ | $V^2_{ijkl} = (\alpha_{ijk} + \beta_{ijk}\hat{X}_{ijkl})^{-1} + \epsilon$ |
| $(iv)$ | $\sigma_{ijkl} = \alpha_{ijk} + \beta_{ijk}\hat{X}_{ijkl} + \epsilon$ |

that the model uses $i$, the mode, $j$ subgrouping, and year as categorical variables in the regression. The similarity of columns $ijk$ and $ij$ indicate that the addition of time as a qualitative variable does not significantly improve the model. For sampling variance estimation, the $R^2$ results in Table 2 demonstrate model $(ii)$ to be superior to the model of Eltinge and Sukasih (2001) and the best fit of the data for the models considered.

Table 2: A table of Adjusted $R$-squared values for models $(i)$ through $(iv)$ using 3-digit grouping of the UCCs

| | $ijk$ | $ij$ | $ik$ | $jk$ | $i$ | $j$ | $k$ |
|---|---|---|---|---|---|---|---|
| $(i)$ | 0.59 | 0.55 | 0.46 | 0.38 | 0.45 | 0.37 | 0.29 |
| $(ii)$ | 0.91 | 0.91 | 0.85 | 0.79 | 0.86 | 0.79 | 0.73 |
| $(iii)$ | 0.87 | 0.84 | 0.79 | 0.61 | 0.77 | 0.60 | 0.49 |
| $(iv)$ | 0.81 | 0.76 | 0.66 | 0.71 | 0.63 | 0.67 | 0.52 |

It should be noted that UCC groupings with a single UCC in the class were not investigated further. A discussion of single UCC variance estimation can be found in Eltinge and Sukasih(2001).

## Hierarchical Bayes Variance Modeling

In this section we attempt to improve our generalized variance function estimate of variance by modeling effects in our variance estimation. The GVF estimator proposed in the previous section is a simple regression model and as such does not account for sampling error. In particular, the regression model assumes that the sample estimate of variance $S^2$ is the true variance $\sigma^2$. We expand the model as a full hierarchical Bayes model to include the error that a design-based variance model misses.

Modeling variance is an area of active research in the field of statistics. Many models estimate the sampling variance of the mean directly from the mean estimates or have assumed the variances to be known. In practice these are not known but estimated through the application of a variance estimation method (e.g. bootstrap, jackknife, balanced repeated replication). These methods do not consider the relationship between groups of variables or the variability of the method used to arrive at the variance estimate. The generalized variance function (see Wolter, 1985) incorporates the relationship between groups of variables but neglects the error introduced by the design-based variance estimator. Arora and Lahiri (1997) modeled these designed based error through the use of a hierarchical Bayes model. This research expands that model to include the generalized variance function within the hierarchical model.

We propose a hierarchical Bayesian model of the form

Level 1

$$\frac{\nu S^2_{ijkl}}{\sigma^2_{ijkl}}|\sigma^2_{ijkl} \sim \chi^2_\nu$$

Level 2

$$\log \sigma^2_{ijkl}|\alpha_{ij}, \beta_{ij} \sim N(\alpha_{ij} + \beta_{ij} \log \hat{X}_{ijkl}, \sigma^2_e)$$

Level 3

$$\alpha_{ij} \sim N(0, 10000)$$

$$\beta_{ij} \sim N(0, 10000)$$

$$\sigma_e^{-2} \sim Gamma(0.001, 0.001)$$

to take advantage of the linear relationship between the groups of variables and estimate the sampling variance.

Initial Model Description

When developing a Bayesian model we must always consider the exchangeability of variables. Parameters are exchangeable when their joint distribution is invariant to permutations in their indexes. When applied to consumer expenditure variance, we want to say that a UCC's variance is exchangeable with any other UCC variance regardless of year, group, and survey mode. Clearly, some indexes do contain information and should not be considered exchangeable. When we include the relevant information as covariates in the model we have the exchangeability that is desired.

The model proposed is a hierarchical model that uses exchangeability at different levels of the model. At the top level we say that $\frac{\nu S^2_{ijkl}}{\sigma^2_{ijkl}}|\sigma^2_{ijkl} \overset{iid}{\sim} \chi^2_\nu$. This would say that the ratio of variances has a $\chi^2$ distribution and that there is a common degree of freedom, $\nu$, for all indexes. Since the sample variances are balanced repeated replicate variance estimates from 44 replicates we have 43 degrees of freedom. This degree of freedom assumption will be tested with a second model.

At the next stage in the model, we apply the generalized variance model developed in the previous section. Here we say that $\log \sigma^2_{ijkl}|\alpha_{ij}, \beta_{ij}$ has a Gaussian distribution with the mean, $\alpha_{ij} + \beta_{ij} \log \hat{X}_{ijkl}$, from the GVF estimate and common error term, $\sigma^2_e$. When the impact of mode, $i$, and group $j$ are accounted for in the regression coefficients the remaining effects are exchangeable Gaussian affects with common variance.

At the final stage we model hyperparameters in our regression model. In theory, we would like to have a flat prior on each of the hyperparameters. However, using a flat prior could lead to an improper posterior distribution. In practice, a diffuse proper prior is used so that the resulting posterior distribution is also proper. A diffuse proper Gaussian prior was used to describe the form of the regression coefficients. This effectively flat prior causes the posterior to be proper without imposing any bias mean value of the coefficients. Similarly, at this stage a diffuse proper prior is placed on the variance. The proper prior for variance is currently under discussion in literature. Some feel that a uniform distribution on the nonnegative real values is a less informative prior than an inverse gamma distribution in some variance applications. This research used an inverse gamma prior for the variance hyperparameter. Posterior distribution estimates for each of the regression coefficients, the variance estimators, and the regression error term can be calculated using Markov Chain Monte Carlo techniques.

Implementation

The high dimensionality of the model lends itself to the use of Gibbs sampling. The model was implemented using WinBUGS and with an R interface. The full model has three groups of UCCs with a total of 60 UCCs in the three groups. Applied to the data sets five modes and the three classification groups leads to thirty regression parameters, 3300 variance parameters and one variance parameter from the second level of the model. The number of parameter in the full model is quite large so the model was implemented in each group separately. Some loss of information in $\sigma^2_e$ will result from the reduced model but with large data set this loss is minimal.

An initial run was slow to converge. Gelman and others proposed that collinearity between parameter is usually the cause of slow convergence (Gelman et. al. (1995), WinBUGS Manual). Collinearity between the regression coefficients was identified and reduced by changing the parameterization of the mean by simply centering the logarithm of the estimates $\log \hat{X}$ on their mean. A second run with a sample of 2000 replicates was drawn from the posterior distribution in three chains after a burn-in period of 7500 draws. The chains were seeded with random values from their prior distributions. Values of the random variables in the three chains were monitored to insure the burn-in was adequate to remove the affect of initial values in the draws. The modified Gelman-Rubin statistic, $\hat{R}$, was also monitored for convergence of the chains.

In WinBUGS, variables are represented by nodes that are either deterministic or stochastic. This imposed a

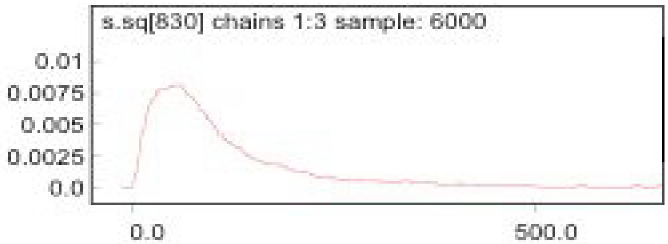


Figure 1: Posterior predictive distribution plot for UCC 280120 for mode 4 in year 1988. Bedroom Linens

problem for the first stage of the model in that the ratio on the left hand side is a deterministic quantity but that ratio has a $\chi^2$ distribution and is therefore considered stochastic, too. As it is currently written, WinBUGS will not allow the ratio to be both deterministic and stochastic. WinBUGS does allow us to implement our own distributions through the likelihood function of the random variable. The first level was rewritten as $S^2_{ijkl}|\sigma^2_{ijkl} \overset{iid}{\sim} \sigma^2_{ijkl}\chi^2_\nu$ and was implemented using the zeros technique which generates a posterior draw from the likelihood function of $S^2_{ijkl}$ through the Poisson distribution (WinBUGS Manual).

Initial model results

Table 3 is a subset of the and results when implemented in WinBUGS. Given are the posterior distributions bounds for the 95% credible set, the posterior mean and the standard deviation for 6000 samples after a burn in period of 7500 draws. Note in the table that alpha[i] is the constant term in the GVF model, beta[i] is the slope parameter and sigma is the variance from the Gaussian distribution in level 2 of the model. Also included in the table is a set of draws from the posterior predictive distribution for bedroom linens.

The 95% credible sets for $\alpha_i$ show that the constant parameter shows a distinct difference between the interview source constants, alpha[1]- alpha[3] and the diary constants, alpha[4] and alpha[5]. The beta credible sets indicate that there is little difference in the slope parameter.

The sample of 6000 draws for each parameter allows us to plot the posterior distribution of the mean expenditure. A plot of the posterior distribution for the sampling variance for the mean expenditure showed the mostly Gaussian form of the posterior distribution. The posterior predictive distribution given in figure 1 for the a UCC, mode and year shows the $\chi^2$ form of $\sigma^2$.

Table 3: WinBUGS output for regression coefficients, common error term, and posterior distribution statistics for UCC 280120, bedroom linens, for all five survey modes in 1988

| node | mean | sd | 2.5% | 97.5% |
|---|---|---|---|---|
| alpha[1] | 0.45 | 0.07 | 0.32 | 0.57 |
| alpha[2] | 0.62 | 0.07 | 0.48 | 0.75 |
| alpha[3] | 0.74 | 0.07 | 0.59 | 0.88 |
| alpha[4] | 1.83 | 0.07 | 1.69 | 1.97 |
| alpha[5] | 1.84 | 0.07 | 1.70 | 1.98 |
| beta[1] | 1.68 | 0.07 | 1.55 | 1.81 |
| beta[2] | 1.61 | 0.07 | 1.48 | 1.74 |
| beta[3] | 1.67 | 0.07 | 1.53 | 1.81 |
| beta[4] | 1.62 | 0.07 | 1.49 | 1.74 |
| beta[5] | 1.80 | 0.07 | 1.67 | 1.93 |
| sigma | 0.84 | 0.02 | 0.80 | 0.88 |
| sigma.sq[111] | 0.08 | 0.02 | 0.05 | 0.13 |
| sigma.sq[112] | 0.11 | 0.02 | 0.07 | 0.16 |
| sigma.sq[113] | 0.15 | 0.03 | 0.10 | 0.22 |
| sigma.sq[114] | 0.99 | 0.23 | 0.63 | 1.52 |
| sigma.sq[115] | 0.69 | 0.15 | 0.44 | 1.05 |

Convergence Diagnostics

Convergence techniques are continually being developed as models with exceptions to convergence are found. In practice a number of convergence techniques should be implemented to verify convergence. To determine if the number of draws is adequate we look at the Monte Carlo error. We desire it to be less than 0.05. For our analysis the MC error was less than 0.001. The number of iterations is reasonable if the model has converged on the posterior distribution.

To analyze convergence further three chains with initial values that were overdispersed were run and a plot of the variables current values were created. These indicated that the chains are well mixed signifying convergence. Lastly, the Gelman-Rubin plot, which is an ANOVA like measure plotted in intervals of 50 samples, demonstrated convergence of the chains to the posterior.

Other Models

Under the current model the degree of freedom parameter in the model is fixed for all $\sigma^2$ parameters at $\nu = 43$. Preliminary empirical analysis showed that the degrees of freedom could be significantly less than 43. A second model with $\nu = 5$ was implemented to see if there was improvement in modeling the variance. Once again, the degree of freedom is constant for all $\sigma^2$ and not allowed to vary over collection mode or UCC grouping.

A third more naive model, below, removes one level of hierarchy in the model. This model allows for the GVF es-

Table 4: Divergence Measures for three models

| UCC | $\nu = 43$ | $\nu = 5$ | naive model |
|---|---|---|---|
| bedroom linens | 11868 | 11960 | 10049 |
| nonelectric cookware | 324 | 233 | 236 |
| power tools | 2640 | 2430 | 2548 |

timate but no longer imposes that the variance come from a $\chi^2_\nu$ distribution. Since we do not have the true variance but only a BRR estimate of variance we can expect that this will reduce the variance of the variance estimate. This model does allow for model error through the precision of the log-normal distribution.

Naive model

$$S^2_{ijkl}|\alpha_{ij}, \beta_{ij} \sim \log N(\alpha_{ij} + \beta_{ij} \log \hat{X}_{ijkl}, \sigma^2_e)$$
$$\alpha_{ij} \sim N(0, 10000)$$
$$\beta_{ij} \sim N(0, 10000)$$
$$\sigma_e^{-2} \sim gamma(0.001, 0.001)$$

Model Comparison

The models were compared using simulated data from the posterior predictive distribution and the observed values. The method of Laud and Ibrahim (1995) uses a divergence approach to model selection. The observed data vector, $y_{obs}$ in compared to the predicted data vectors $y^{(l)}_{new}$. The smaller the divergence the better the model fit to the data. The divergence $d(y_{new}, y_{obs}) = E\{n^{-1}||y_{new} - y_{obs}||^2|y_{obs}\}$ is approximated by $(nB)^{-1}\sum_{l=1}^{B}||y^{(l)}_{new} - y_{obs}||^2$, where $B$ is the number of replications per variable and $n$ is the number of variables replicated. One problem discovered with the divergence measure of Laud and Ibrahim is that the divergent measure is not unitless.

The results of the analysis are shown in Table 4 above. Note that no model is uniformly superior over the UCCs considered. This is true in general for all 60 UCCs. All of the models have divergence measure in the same order of magnitude and therefore no model is unilaterally rejected or accepted based on the Laud and Ibrahim divergence measure.

## Conclusions

One explanation for the poor performance of the large $\nu$ model is that the variance estimates from the GVF regression are misspecified by the $\chi^2_\nu$ distribution. The constant degree of freedom for all UCCs may be to restrictive causing undue variance in the variance estimate of some UCCs. One way to compensate for this shortfall is to model the degree of freedom. This will be investigated in future research.

Further exploration into variance modeling is necessary before any conclusive results can be garnered from this research. Of particular interest in this area is modeling the variance matrix rather than the univariate variance. Since this research used WinBugs we were limited to the univariate case.

One direct outcome of value is a variance estimate of the BRR variance. This value could quantify when a statistic using BRR methods is publishable. A measure of consistency of variance has been achieved through this research.

## References

Arora, V. and Lahiri, P. (1997), 'On the superiority of the bayesian method over the blup in small area estimation problems', *Statistica Sinca* **7**, 1053–1063.

Bell, W. (1998), Accounting for uncertainty about variances in small area estimation, Technical report, U.S. Bureau of the Census.

Congdon, P. (2001), *Bayesian Statistical Modelling*, 1 edn, Wiley.

Data, G. S., Lahiri, P., Maiti, T. and Lu, K. L. (1999), 'Hierarchical bayes estimation of unemployment rates for the states of the u. s.', *Journal of the American Statistical Association* **94**(448), 1074–1082.

Eltinge, J. L. and Sukasih, A. (2001), Approximation methods for covariance matrix estimators usen in analysis of diary and interview data from the u.s. consumer expenditure survey, BLS document.

Eltinge, J. L., Sukasih, A. and Weber, W. (2000), Feasibility of constructing combined estimators using consumer expenditure interview and diary data, Paper presented to the Bureau of Labor Statistics Conference on Issues in measuring Price Change and Consumption.

Fay, R. E. and Herriot, R. A. (1979), 'Estimates of income forr small places: An application of james-stein procedures to census data', *Journal of the American Statistical Association* **74**(366), 269–277.

Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. (1995), *Bayesian Data Analysis*, 1 edn, Chapman Hall.

Gill, J. (2002), *Bayesian methods for the social and behavioral sciences*, Chapman and Hall.

Johnson, E. G. and King, B. F. (1987), 'Generalized variance functions for a complex sample survey', *Journal of Official Statistics* **3**(3), 235–250.

Laud, P. and Ibrahim, J. (1995), 'Predictive model selection', *Journal of the Royal Statistical Society* **57**(1), 247–262.

Neter, J., Kutner, M. H., Nachtsheim, C. J. and Wasserman, W. (1996), *Applied Linear Statistical Models*, 4 edn, WCB/McGraw-Hill.

Otto, M. and Bell, W. (1995), 'Sampling error modelling of poverty and income statistics for states', *American Statistical Association, Proceedings of the Section on Government Statistics* pp. 160–165.

U.S (1997), *BLS Handbook of Methods*, chapter 16: consumer expenditures and income edn. U.S. Government Printing Office.

Valliant, R. (1987), 'Generalized variance functions in stratified two-stage sampling', *Journal of the American Statistical Association* **82**, 499–508.

Wolter, K. M. (1985), *Introduction to Variance Estimation*, Springer-Verlag, New York.