A BOOTSTRAP VARIANCE ESTIMATOR FOR SYSTEMATIC PPS SAMPLING

Steven Kaufman, National Center for Education Statistics
Room 422d, 555 New Jersey Ave. N.W., Washington, D.C. 20208

Key Words: Simulation, Half-Sample Replication

1. Introduction

Systematic probability proportionate to size (PPS) sampling procedures (Wolter, 1985, section 7.6) are efficient in terms of ease of selection and lowering sampling error. For this reason they are used extensively in large-scale surveys. Since each stratum systematic sample is selected using a single random start, the sample can be viewed as a sample of size one, where each sample consists of a single sample cluster of n_h primary sampling units (PSUs). Therefore, it is impossible to produce an unbiased variance estimator since the sample size is one. However, a number of biased methodologies are used for variance estimation.

These methodologies generally take one of two forms: 1) assume the systematic sample can be approximated by a simpler sample design with a known variance estimator or 2) assume the response variable follows some super-population model and a variance estimator is produced appropriate for that model. Both these approaches allow for grouping of PSUs, so variances can be computed within groups. Wolter (1985, chapter 7) provides a good discussion of a number of systematic sample variance estimators that can be classified into one of these two forms. An example, using balanced half-sample replication (BHR) is provided below.

BHR is a widely used variance replication methodology for complex survey designs. It is designed for samples where two PSUs within each stratum are selected with replacement. With BHR, choosing one PSU within each stratum generates a half-sample. A number of half-samples are generated by alternating which PSU, within stratum, go into the half-samples. The BHR variance is the simple variance of the half-sample estimates. Through a balancing process of the half-samples, the BHR variance estimate, for linear estimates, equals the direct sample variance estimate.

BHR can be adapted to designs where more than two PSUs are selected in a stratum by consecutively pairing selected PSUs, after placing them in the original order of selection; and assuming each pair is a stratum for variance estimation (variance stratum). If without-replacement sampling is used then a finite population adjustment can be applied. See (Wolter, 1985 pp. 110-152) for a more complete description of BHR.

In order to use BHR with systematic PPS sampling, it must be assumed that a PPS selection can be approximated by the deep stratification induced by the pairing described above. This assumption is reasonable, considering that the first sort variable, ignoring the lack of independence between breaks in the variable, can be considered an implicit stratification. However, BHR also assumes that the variance estimate is proportional to the inverse of the sample size. (This follows from $V_{BHR}(X) = V((X_1 + X_2)/2) = 1/2V(X_1)$, where subscript 1 and 2 represents the estimate based on the first and second PSUs respectively selected in each stratum.) In section 2.0, it will be demonstrated, through a simulation study, that systematic sampling variances are not necessarily inversely proportional to the sample size.

2.0 Using the BHR Model with Systematic Sampling

For the BHR model to work, the stratum variances must be proportional to $1/n_h$, as n_h increases or decreases, since BHR makes this assumption. If this assumption is not true then the BHR model is unlikely to produce accurate results. To investigate the $1/n_h$ assumption, a simulation study is done, using the sample design described in section 4.2. Four thousand systematic PPS samples are selected with sample sizes of n_h and $0.5n_h$. By computing the simple variance of the 4,000 simulation estimate, an estimate of the true variance is computed. This is done for estimates of total students, teachers and schools. If the variance is proportional to $1/n_h$, then the ratio, $R100/50 = \hat{V}_{n_h}(X_l)/\hat{V}_{0.5n_h}(X_l)\times 0.5-1$ should be close to 0; where l represents the estimate type (total students, teachers or schools). When the ratio is less (greater) than 0, the systematic sample variance decreases faster (slower) than the $1/n_h$ assumption would imply. A negative (positive) ratio means that BHR should overestimate (underestimates) the variance.

Table 1 demonstrates that sometimes the ratio is close to 0. Other times, it is a great deal different than 0. The systematic PPS sampling variance does not necessarily decrease faster than the $1/n_h$ assumption would imply; sometimes its decrease is slower. This is an indication that BHR will not necessarily produce an overestimate of the variance, which is a common assumption among sampling statisticians. When there is a large difference from 0, the magnitude is dependent on the variable. This seems to imply, since the sampling rates are not high, that the violation of the $1/n_h$ assumption is due to the initial sort ordering (i.e., within sample correlation).

It should be noted that the table 1 results exaggerate the true impact of the $1/n_h$ assumption. Using the $1/n_h$ assumption, the ratio, used in the table, adjusts the variance with the smaller sample size to approximate the variance with the larger sample size. This approximation uses the smaller sample estimate's unknown finite population correction. Since the true finite population correction is likely larger than the one used in the approximation, the absolute value of the true impact of the $1/n_h$ assumption should be expected to be smaller than what table 1 indicates.

The important conclusion from this example is that variance estimates, based on designs using systematic sampling, will not necessarily be proportional to $1/n_h$, as n_h increases or decreases. When this occurs, an important BHR assumption is violated, and the BHR variance estimator should not be expected to perform well when the magnitude of the violation is large.

The statements concerning the proportionality of the variance estimate are qualified with 'as n_h increases or decreases'. The importance of this qualification can be seen with equal probability systematic sampling. Here, the variance can be expressed proportional to $1/n_h$ (e.g., $V(\bar{y}_h) = [(N_h - n_h)/N_h][S^2_{wst}/n_h][1+(n_h-1)\rho_{wst}]$, see (Cochran, 1977, pp. 209)). If S^2_{wst} and ρ_{wst} are constant for an arbitrary n_h then $V(\bar{y}_h)$ would be approximately proportional to $1/n_h$, as n_h increases or decreases. However, both S^2_{wst} and ρ_{wst} are within systematic sample population estimates. This implies that as n_h changes, the systematic samples change; hence S^2_{wst} and ρ_{wst} also change by some unknown function of n_h. Therefore, even though $V(\bar{y}_h)$ is proportional to $1/n_h$ for fixed n_h, as n_h increases or decreases, the variance may not be proportional or even closely proportional to $1/n_h$.

3.0 Bootstrap Variance Model

To address the situation when the systematic variance in not proportional to $1/n_h$, a bootstrap variance estimator is proposed in this paper, which is less dependent on the $1/n_h$ assumption than the BHR estimator. This section first describes the consistency theorem for the bootstrap estimator; by example, the super-population model, used in the proposed bootstrap procedure, is demonstrated; next, the mechanics of the bootstrap procedure is presented; and finally, the consistency of the bootstrap procedure is established. We begin by describing the super-population model.

3.1 The Consistency Result

Theorem

The required assumptions are:

1) a systematic PPS sample (s_{ih}) has a known partition (i.e., $s_{ih} = \bigcup_{c=1}^{C_{ih}} s_{ihc}$);
2) $\hat{X} = \sum_h \sum_{j\in h} w_{hj} x_j = 1/n \sum_h \sum_{j\in h} y_j$ is the estimate of interest, with w_{hj} being the sampling weight and x_j being the variable of interest;
3) as n increases, the sample allocation between stratum remains constant;
4) for PSUs in s_{ihc}, the y_j's are conditionally i.i.d given s_{ihc} and are generated from an otherwise unspecified distribution function $F_{hc}(y)$ which satisfy conditions for Mallows' distance;

and

5) between partitions, the y_j's are conditionally independent given the s_{ihc}'s, but not identically distributed.

Bullets 1-5 specify the super-population model.

It then follows that the bootstrap variance estimator of $\hat{X}$ given s_{ih} generated from the bootstrap estimates $\hat{X}^*_b = 1/n \sum_h \sum_{j\in h} y^*_j$, where the y^*_j's are generated from $\hat{F}_{hc}(y)$, is consistent, as $n \to \infty$, provided $\hat{F}_{hc}(y) \to F_{hc}(y)$ and $\mu^*_{yhc} \to \mu_{yhc}$, as $n \to \infty$. μ^*_{yhc} is the bootstrap expectation of y within a partition.

The proof follows from the super-population assumptions using the argument in example 3.1 from (Shao and Tu, 1995). The details are provided in (Kaufman, 1998).

3.2 Bootstrap Model Example

In practice, the statistician rarely knows the required partitioning ($s_{ih} = \bigcup_{c=1}^{C_h} s_{ihc}$). However, the statistician usually orders the frame before sample selection. With this ordering, the statistician is implicitly assuming that nearby PSUs are similar, at least in terms of the most important response variables. This implicit assumption can be used to develop a partitioning that approximately meets the required assumptions.

An example is provided below.

For a fixed even numbered sample size (n_h), the elements of the partition (s_{ich}) can be determined by pairing the sets of PSUs within consecutive sampling intervals, after the frame has been placed in its original sort ordering. All samples have the same partitioning (i.e., the partitioning is only a function of stratum, -- s_{ch}, $c = 1$ to C_h) and each s_{ich} (s_{ch}) has exactly two PSUs. In terms of consistency, it is assumed that the partitioning remains fixed as the sample size increases and more PSUs are selected within a partition. This "type" of partitioning is used in the bootstrap procedures proposed in this paper.

An additional observation about this partitioning is:

If the partitioning methodology described above correctly models the distribution of X; the n_h's are even and increase by multiples of C_h then the $\underset{2}{E}(\hat{X}_i | s_h's) = K$, a constant; where $\underset{2}{E}$ refers to the expectation with respect to the super-population model. Therefore,

$$V(\hat{X}_i) = \underset{1}{E}\underset{2}{V}(\hat{X}_i|s_h's) + \underset{1}{V}\underset{2}{E}(\hat{X}_i|s_h's) = \underset{1}{E}\underset{2}{V}(\hat{X}_i|s_h's),$$

where 1 refers to the selection of the $s_h's$. Since the bootstrap variance estimator is consistent for $\underset{2}{V}(\hat{X}_i|s_h's)$, the bootstrap variance is consistent to an unbiased estimator for the unconditional variance.

3.3 Bootstrap Sample Size (n_h^*)

Since it is assumed that the relationship between the variance and n_h is unknown, the actual bootstrap sample size (n_h^*) used in the bootstrap selections must be computed through a series of trial and error simulations. This is done by comparing and estimate of the true variance with the bootstrap variance for a specific bootstrap sample size.

Determining n_h^* through a simulation provides a robust variance estimate because $V^*(\hat{X}_h)$, by construction, will be almost unbiased, even if the model assumptions are false. The disadvantage of the simulation is that it can only be implemented with frame variables. However, if n_h^* is relatively flat for non-frame variables, the bootstrap replicate weights should be applicable for those variables, too.

3.4 Bootstrap Implementation

To perform the simulation study, frame variables are used, so estimates can be computed for any selected sample. The statistician always has three estimates available for this purpose. One is the measure of size or some function of the measure of size. The second is the estimate of the total number of PSUs (sum of the sample weights). The third is the average measure of size per PSU or the average per PSU of some function of the measure of size. If the measure of size is used in the simulation, it will be necessary to use a different year's data to produce estimates; otherwise, the variances will be zero.

To determine the appropriate n_h^*'s, the simulations must first be applied to individual stratum estimates Θ_h. The simulation process for estimating the bootstrap variance, $V^*(\Theta_h)$ for an estimator Θ_h, works as follows:

3.4.1 Bootstrap Procedures

1. Select a sample (s_i) from the original frame, using the PPS methodology of the original sample design.
2. For the initial bootstrap sample size values, n_h^*, use n_h. After the initial simulation, n_h^* will likely require adjustment for at least some of the strata.
3. Generate a bootstrap frame based on the selected sample. For each selected PSU j, w_j bootstrap PSUs (bj) are generated by replicating the j^{th} PSU w_j times. The bj^{th} bootstrap-PSU has the following measure of size (m_{bj}):

$$m_{bj} = I_{bj} \cdot 1/w_j,$$

$$I_{bj} = \begin{cases} 1, \text{ if } bj \text{ is an integer component of } w_j \\ C_i, \text{ if } bj \text{ is a noninteger component of } w_j \\ \quad C_j \text{ being the noninteger component} \end{cases}$$

4. Randomize the bootstrap frame according to super-population model specification. This is accomplished by placing the bj bootstrap-PSUs generated from PSU j within stratum h and sample s_i in their original order of selection. Next, bootstrap-PSUs generated from the first PSU are paired with the next set of bootstrap-PSUs generated from the second PSU. The third set of bootstrap-PSUs is paired with the fourth set. This process continues until all bootstrap-PSUs are paired. If there are an odd number of PSUs then the last set of groupings of bootstrap-PSUs contains the bootstrap-PSUs generated from the last three PSUs in stratum h. This is repeated for every stratum in s_i. Now, the bootstrap-PSUs are randomized within their respective pair.
5. The bootstrap frame, bootstrap frame ordering, measure of size (m_{bj}), and bootstrap sample size (n_h^*) have been specified. Select B bootstrap samples, after re-randomizing the bootstrap-PSUs

after each selection, using the same procedures used to select the original systematic PPS sample. The one exception to this is that a bootstrap-PSU generated from noncertainty PSUs that become certainty in the bootstrap selection should not be eliminated from the selection process and taken in sample with probability 1. Their selection probability should remain unchanged and if the bootstrap-PSU is selected multiple times that should be reflected in the bootstrap weight (see 6 below).

6. For each bootstrap sample, b, compute a set of bootstrap weights, w_j^{*b}, and then compute Θ_{ibh}^* by using w_j^{*b} instead of w_j in the formula for Θ_h.

   The bootstrap-PSU weight, w_j^{*b}, is:

   $w_j^{*b} = \sum_{bj \in S_j^b} w_{bj}^p$, S_j^b is the set of all bj generated from j that are selected in the b^{th} bootstrap sample.

   and $w_{bj}^p = I_{bj} \cdot M_{bj} / p_{bj}$

   M_{bj}: is the number of times the bj^{th} bootstrap- PSU is selected,

   p_{bj}: is the bootstrap selection probability for the bj^{th} bootstrap-PSU.

   $p_{bj} = m_{bj} / SI_h, SI_h = \sum_{bj \in s_h} m_{bj} / n_h^*$.

7. The bootstrap variance is:

   $V^*(\Theta_{ih}) = 1/(B-1)\sum_{b=1}^{B}(\Theta_{ibh}^* - \overline{\Theta}_{ih}^*)^2$,

8. Repeat steps 1-7, for a large number of samples (T).
9. Compute the simple variance of Θ_{ih} from $i = 1$ to T, $\hat{V}(\Theta_h)$, as a measure of the true variance; and compute the average bootstrap variance $\overline{V}^*(\Theta_h)$, averaged over the T, $V^*(\Theta_{ih})$ estimates.
10. Compare $\overline{V}^*(\Theta_h)$ with $\hat{V}(\Theta_h)$ and adjust n_h^* to reduce the bias between $\overline{V}^*(\Theta_h)$ and $\hat{V}(\Theta_h)$.
11. Repeat steps 1-10, until this bias has been reduced to a satisfactory level.
12. Using the n_h^* from step 11, repeat steps 3-6 for the actual collected sample, generating a set of bootstrap replicate weights, w_j^{*b} that can be used to compute variances of other, more complex statistics that are not necessarily computed within h.

3.5 Consistency of the Bootstrap Estimator

$\hat{F}_{hc}(y) \to F_{hc}(y)$ and $\mu_{yhc}^* \to \mu_{yhc}$, as $n \to \infty$ follows from $E_*(\sum_{bj \in D} w_{bj}^p X_{bj}) = \sum_{bj \in D} I_{bj} X_{bj} = \sum_{j \in D} w_j X_j$, where D is a domain and E_* is the bootstrap expectation. See (Kaufman, 98) for details.

4.0 Simulation

To demonstrate the advantages of the bootstrap variance estimator, a simulation study is presented comparing BHR and the bootstrap variance estimator. Two thousand simulations, denoted by s, are generated using frame variables. The frame is the National Center for Education Statistics' (NCES) Private School Survey (PSS). The PSS is NCES's school frame for private elementary and secondary schools. Three totals (number of schools, number of teachers, and number of students), two averages (average students and average teachers per school), and one ratio (ratio of number of students to number of teachers) are estimated in the simulation. In tables 3-5, estimates are computed by each stratification variable (affiliation, region and school level), as well as one of the sort variables (Urbanicity). The School and Staffing Survey (SASS) sample design is used to select the simulation samples. Relative error, relative mean square error, and coverage rates are used to measure performance.

4.1 Comparison Statistics

In this section, the statistics used to compare the bootstrap and BHR variances are described.

4.1.1 Relative Error

Rel. Error $= (\overline{V}_e(\Theta)^{1/2} / V_t(\Theta)^{1/2} - 1) \cdot 100$

Where: $\overline{V}_e(\Theta)$ is the average of the variance estimates ($V_e(\Theta_s)$) from either the bootstrap or BHR procedure.

$V_t(\Theta) = 1/1999 \sum_{s=1}^{2,000} (\Theta_s - \overline{\Theta})^2$.

4.1.2 Relative Mean Square Error

$\{[VV_e(\Theta) + (\overline{V}_e(\Theta) - V_t(\Theta))^2]^{1/2} / V_t(\Theta)\} \cdot 100$,

Where: $VV_e(\Theta) = 1/1999 \sum_{s=1}^{2,000} (V_e(\Theta_s) - \overline{V}_e(\Theta))^2$.

4.1.3 Coverage Rates

The coverage rate is the percent of the time the 95% confidence interval contains the true value.

4.2 SASS Sample Design

The sample frame, used in the simulation, is the list frame component of NCES's Private School Survey (PSS). The list frame is stratified by detailed School Association (19 groups), within Association by Census Region (4 levels), and within Region by school level (elementary, secondary and combined). The school sample is selected using the systematic probability proportionate to size sampling procedure, described in the introduction. The measure of size is square root of the number of teachers in the school. Before sample selection, the school frame is ordered by state, school highest grade, urbanicity, zip code, and school enrollment. To reduce the time to complete 2,000 simulation only one detailed school association is used.

4.3 Determining n_h^* for the Bootstrap Variance

As described in section 3.3, the determination of n_h^* requires a simulation study in itself. For each stratum, a series of simulations was done for various n_h^*. The optimum n_h^* is likely dependent on the estimate of interest. Since we want only one set of replicate weights, a compromise n_h^* is determined that works reasonably well for all estimates. The results presented below use the compromise set of n_h^*. Table 2 presents the values for n_h and n_h^*. Each simulation used in the determination of n_h^* had at least 250 samples.

4.4 BHR Variances

The r^{th} school half-sample replicate is formed using the usual textbook methodology (Wolter, 1985) for establishment surveys with more than 2 units per stratum. This is described in the introduction. Two BHR variance estimates are presented. The first (BHR without FPC Adjustment) is the variance estimates described above. This estimate does not make any type of Finite Population Correction (FPC) adjustments. The second BHR variance estimate (BHR with FPC Adjustment) adjusts the first variance estimator by $1 - P_h$, where P_h is the average of the selection probabilities for the selected units within stratum h.

4.5 Number of Replicates

Thirty-two and thirty replicates have been used in the BHR and bootstrap variances, respectively.

4.6 Results

Because of space consideration, three tables have been excluded from this paper. These tables are included in (Kaufman, 1998).

According to tables 3-5, in terms of extremes, the bootstrap variance estimator is better than either BHR variance estimator with respect to relative error, relative MSE, or coverage rate. The bootstrap relative errors are large in absolute value (greater than 20% or less than – 20%) once, while the BHR, with and without FPC adjustment, relative errors are large 8 and 4 times, respectively.

Only 5 of the bootstrap relative MSEs are larger than 50% and none are greater than 100%. The BHR without FPC adjustment has 18 relative MSEs larger than 50% and 3 greater than 100%. The FPC adjusted BHR has 14 relative MSEs larger than 50% and 2 larger than 100%.

The bootstrap procedure has no high coverage rates (coverage rate greater than 98%) and 1 low coverage rates (coverage rate less than 89%). The bootstrap has no coverage rate greater than 99%. The BHR without FPC adjustment has 7 high coverage rates, no low coverage rate and 5 larger than 99%. Even with a FPC adjustment, the BHR has 6 high coverage rates, 1 low coverage rate, and 5 coverage rates greater than 99%.

The difference between the bootstrap and BHR is largest for the Urbanicity estimates. For these estimates the BHR relative MSE can be almost 4 times larger than the bootstrap relative error (see tables 3 and 4 Urban). One possible explanation for this may be that the Urbanicity sample size is indirectly controlled by the third sort variable, while the other estimates are directly controlled by the stratification.

4.7 Conclusion

This paper discussed how BHR can be used to measure the variances from surveys utilizing systematic PPS selection procedures. Two assumptions are necessary: 1) the extra stratification introduced by the variance stratum is sufficient to reflect the systematic process and 2) the variance is inversely proportional to the sample size. In table 1, it has been observed that systematic PPS sampling variances may not be inversely proportional to the sample size.

To correct this problem, a bootstrap variance estimator has been introduced which does not make the inverse sample size assumption. Given an appropriate super-population model, the bootstrap procedure produces consistent variance estimates. Based on the simulation of the SASS survey design (Tables 3-5), the bootstrap variance estimator performs better than the BHR with respect to relative error, relative MSE and coverage rates. This is especially true with the Urbanicity estimates. One drawback of the proposed bootstrap procedure is that the determination of an appropriate bootstrap sample size can only be implemented using frame variables. However, with appropriate frame variables, the bootstrap variances are close to unbiased, even when the super-population model assumption fails.

5.0 References

Cochran, W. (1977) *Sampling Techniques*. New York: John Wiley and Sons.

Kaufman, S. (1998), *A Bootstrap Variance Estimator for Systematic PPS Sampling*, Working Paper No. 98-12, U.S. Dept. of Education, NCES, Washington, D.C.

Shao S. and Tu, D. (1995). *The Jackknife and Bootstrap*. New York: Springer-Verlag.

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

Tables 1 – Measurement of degree the true systematic sampling variance is proportional to $1/n_h$ with respect to different sample sizes

Stratum (h)	n_h/N_h (%)	Teacher R100/50 (%)	Student R100/50 (%)	Schools R100/50 (%)
01911	2.0	-31.2	14.1	-28.0
01912	2.8	-27.0	-2.5	-14.3
01914	3.4	-23.8	-5.3	-19.3
01931	4.5	23.7	18.0	4.6
01932	4.9	2.4	-25.8	3.9
01934	4.3	-20.4	-26.3	-7.0

Table 2 – Original (n_h) and Bootstrap (n_h^*) Sample Size by Stratum

Stratum	n_h	n_h^*	Stratum	n_h	n_h^*	Stratum	n_h	n_h^*
01911	14	12	01921	10	5	01931	48	35
01912	16	11	01922	10	8	01932	46	33
01913	52	28	01923	10	10	01933	114	81
01914	34	24	01924	10	10	01934	52	40

Table 3 -- % Relative Error, % relative Mean Square Error and % coverage rates for the Bootstrap and BHR variance estimator for Schools estimates by Affiliation, Region, Level and Urbanicity

Estimate	Bootstrap			BHR without FPC Adjustment			BHR with FPC Adjustment		
	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate
Other Affil.	-4.5	27.8	93.1	12.7	40.7	97.0	9.7	35.4	97.0
Northeast	4.3	43.6	94.6	10.3	52.3	94.9	8.0	49.0	94.9
Midwest	4.2	42.8	92.9	12.5	51.4	98.3	9.8	46.9	95.7
South	-10.9	32.7	90.7	-6.6	26.5	89.6	-10.3	29.0	89.4
West	-2.4	35.1	92.9	7.8	43.7	92.2	5.1	40.0	92.2
Elementary	1.3	34.9	93.6	16.1	57.0	95.9	14.0	52.9	95.9
Secondary	-2.9	57.0	90.5	26.3	107.1	97.2	14.8	81.9	95.9
Combined	-6.2	29.5	91.2	-1.1	28.2	92.3	-4.2	27.7	92.3
Rural	7.5	36.8	95.7	24.2	71.2	98.7	20.9	63.8	98.7
Suburban	6.5	36.6	95.0	23.1	67.5	97.4	19.9	60.6	97.4
Urban	11.5	43.2	96.1	53.7	147.6	97.5	49.5	135	97.5

Table 4 -- % Relative Error, % relative Mean Square Error and % coverage rates for the Bootstrap and BHR variance estimator for Teachers per School estimates by Affiliation, Region, Level and Urbanicity

Estimate	Bootstrap			BHR without FPC Adjustment			BHR with FPC Adjustment		
	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate
Other Affil.	-5.8	28.4	92.4	4.0	27.3	95.9	1.0	24.6	95.8
Northeast	2.1	42.1	93.7	0.6	41.2	90.9	-1.7	39.4	90.8
Midwest	-0.7	37.7	92.5	18.2	60.7	99.7	15.0	54.0	99.7
South	-10.9	32.5	89.4	-9.6	28.0	89.4	-13.2	31.4	88.1
West	5.5	41.3	95.1	12.1	45.4	93.6	9.2	40.1	93.6
Elementary	4.6	38.7	94.0	17.7	57.2	97.1	15.3	52.4	97.1
Secondary	8.6	54.2	95.2	29.4	93.3	97.4	16.6	63.8	93.7
Combined	-6.9	29.9	91.6	-4.1	26.2	92.3	-7.2	27.1	91.0
Rural	1.1	37.4	93.2	27.9	83.6	99.6	24.2	74.6	99.6
Suburban	-10.7	34.4	89.7	-2.9	34.9	91.8	-5.5	34.3	91.8
Urban	10.6	44.9	95.5	61.6	177.	99.8	56.7	161.4	99.8

Table 5 -- % Relative Error, % relative Mean Square Error and % coverage rates for the Bootstrap and BHR variance estimator for Students/Teacher Ratio estimates by Affiliation, Region, Level and Urbanicity

Estimate	Bootstrap			BHR without FPC Adjustment			BHR with FPC Adjustment		
	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate	Rel. Error	Rel. MSE	Cov. Rate
Other Affil.	-0.3	31.9	94.2	12.3	46.0	96.1	9.1	40.6	94.7
Northeast	-4.3	56.3	91.6	9.3	70.4	94.7	6.7	66.4	94.7
Midwest	-5.5	67.1	91.2	3.1	43.9	95.4	0.5	41.6	94.1
South	6.7	45.2	95.4	3.2	31.5	95.9	-0.8	28.6	93.4
West	-1.0	38.3	93.9	9.8	46.9	97.3	7.1	42.8	97.3
Elementary	-2.5	43.9	93.2	11.8	52.2	99.5	9.5	48.5	99.5
Secondary	-25.3	49.1	81.1	1.2	33.2	94.1	-9.3	32.5	91.4
Combined	9.7	46.2	95.8	16.3	53.0	95.9	12.6	45.7	95.7
Rural	7.0	59.8	95.3	22.7	84.2	99.9	18.9	75.6	99.8
Suburban	1.1	37.6	93.5	18.0	61.6	97.1	14.8	55.3	97.1
Urban	5.4	45.6	94.7	15.0	58.3	93.8	11.6	52.1	93.7