A New Model for Estimating the Variance under Systematic Sampling

Steven Kaufman, National Center for Education Statistics
room 9059,1990 K St. NW, Washington, D.C 2006

Key Words: Bootstrap, PPS Sampling, BHR, Simulation

1.0 Introduction

Systematic sampling (either with equal or unequal selection probabilities) is a common sampling scheme within statistical organization. It is used because of its simplicity of implementation and its potential increase in efficiency.

A probability proportionate to size (PPS) systematic sample of size n_h can be selected within stratum h, by partitioning the units of a fixed-ordered frame ($\mathbf{F}_{1h}$) into n_h zones, such that the sum of the individual measures of sizes (MOS) in each zone is equal across all the zones. The total MOS within each zone is called the sampling interval, I_h (i.e., the sum of unit MOS in the stratum divided by n_h). A uniform random number between 0 and the sampling interval is chosen, say r_h. Starting with the first unit in the ordered stratum, cumulate the unit MOS. Each unit is assigned a cumulative MOS, C_{hi}, equal to its MOS plus the cumulative MOS of all previously cumulated units. The first unit with cumulative MOS greater than or equal to r_h is the first selected unit. Subsequent selections correspond to the first respective units with cumulative MOS greater than or equal to $r_h + I_h, \ldots, r_h + (n_h - 1)I_h$. Note that one unit is selected within each zone and that each zone acts as an "implicit stratum" (g), in that each zone represents a "subpopulation". The "implicit strata" are not (real) strata because the selections are not independent and because some units are in multiple implicit strata (i.e., some units have a positive selection probability in multiple implicit strata). When the frame is ordered so the units within an "implicit stratum" are relatively homogeneous, one expects the systematic sample, as in a stratified sample, to be more efficient than a completely random selection within the stratum.

One problem with systematic sampling is that such samples can be viewed as a cluster sample of cluster sample size one. As such, unbiased variance estimation becomes impossible without additional assumptions. One common method for approximating the variance from systematic sampling is to treat the implicit strata as real strata when computing variances. Within a sampling stratum, this is accomplished by placing the sample or frame into the original frame ordering before sample selection and consecutively pairing the sample units or implicit strata. Each pair can then be treated as a stratum (variance-stratum) for the purpose of variance computation. Since many sampling schemes have unbiased variances estimators with two units per stratum, the statistician can choose the one most appropriate for the particular sample design and use it to approximate the systematic variance-stratum (vs) variances.

The basic assumption is that most of the efficiency of the systematic sample comes from the implicit stratification. An additional assumption is that the MOS are fixed, non-random quantities known for all units on $\mathbf{F}_{1h}$. In this setting, a sample selected from $\mathbf{F}_{1h}$ will be termed a fixed-ordered sample design.

There are two main concerns with the fixed-ordered variance-stratum approach. The first is that the variance-stratum variances still do not reflect the appropriate systematic sampling variance. As such, the variance may only reflect with-replacement sampling. If the sampling rates are high, it becomes difficult to determine an appropriate finite population adjustment (FPC). An apparent conservative approach would be to apply no FPC. However, it is possible for an appropriate FPC to be greater than 1 (see (2) below). In which case, applying no FPC may not be conservative. Without knowing the correct variance-stratum variance, the variance will be in error even if the variance-stratum assumption is correct (e.g., there may be a nonzero within variance-stratum correlation). The second concern is the variance-stratum assumption. Namely, the correlation between variance-strata may not be zero. Since systematic sampling is considered efficient, it seems like these concerns should lead to an overestimate of the variance. However, since the correlations can be positive and an appropriate FPC can be greater than 1, this need not be the case.

The idea behind the proposed model is to relax the fixed-ordered sample design assumption. Instead, it will be assumed that the MOS are "locally-random". With this assumption, along with the assumption that the correlation between the i^{th} and j^{th} selected PSUs is negligible, it becomes possible to derive an appropriate FPC for a PPS systematic selection procedure, provided that the frame is ordered in some fashion by the MOS before sample selection. This will be presented in the following sections.

2.0 Estimates from a Systematic Sample

PPS systematic sampling is a common procedure used with complex sample designs. One way of selecting such a sample was provided in section 1.0. It will be assumed that: 1) a single stage PPS sample is selected; 2) primary sample units (PSUs) with MOS larger than the sampling interval are added to the sample with certainty and excluded from the sampling procedures; 3) n_h is even for every stratum; and 4) before sample selection, the frame is ordered in some fashion by the MOS.

2.1 Estimating a Total $\hat{T}_{sy}$

An unbiased estimate, $\hat{T}_{sy}$, for the total is $\sum_{h=1}^{H}\sum_{i=1}^{n_h} x_i / p_i$, where H is the number of sampling strata, n_h is the number of sampled PSUs in sampling stratum h, x_i is the value for some variable for selected PSU i, and p_i is the selection probability for the PSU (i.e., p_i is the MOS for PSU i divided by the stratum sampling interval).

2.2 Estimating the Variance, $V(\hat{T}_{sy})$, of $\hat{T}_{sy}$,

$V(\hat{T}_{sy})$ can be express as:

$$\sum_{h=1}^{H}\left[\sum_{i=1}^{n_h} V(x_i / p_i) + \sum_{i=1}^{n_h}\sum_{j\neq i}^{n_h} \rho_{hij}\sqrt{V(x_i / p_i)V(x_j / p_j)}\right] \quad (1)$$

where: ρ_{hij} is the correlation between x_i / p_i and x_j / p_j (i.e., the correlation between the i^{th} and j^{th} PSUs selected in the systematic selection process).

Of course, without further assumptions, none of the above quantities have unbiased estimates.

2.3 "Locally-random" Assumption

The MOS m_i and m_j for PSUs i and j are "locally-random", if there exists a partitioning of the frame, denoted by P_k, such that i and $j \in P_k$ imply m_i and m_j are generated from some random distribution with mean $\mu_i = \mu_j = \mu_k$ and variance $\sigma_i^2 = \sigma_j^2 = \sigma_k^2$. Assuming PSUs are ordered in some way by m_i, before sample selection, the "locally-random" assumption means that PSUs within P_k can be considered to be in a random order. This randomization can be used to determine a variance estimator with an appropriate FPC for a systematic PPS sample design.

There are a number of ways to justify the "locally-random" assumption. In most survey measurements, y_i, it is expected that there exists some response variance. This can be represented by assuming $y_{i\alpha} = \mu_i + e_\alpha$, where μ_i is a constant; and e_α is a random variable with mean zero and variance σ_i^2 which represents the error measurement from the α^{th} repeated measurement. Since m_i comes from or is based on some sort of frame collection, there is no reason to expect m_i to be free of measurement error. So the model $y_{i\alpha} = \mu_i + e_\alpha$ would seem appropriate. This justifies the random assumption. The equality of the μ_i and σ_i within each partition is just a local homogeneity assumption, which is often make when estimating variances.

The "locally-random" assumption can also be justified in term of a super-population model, instead of a response error model.

2.4 Defining P_k

The first step in determining P_k is to place $\mathbf{F}_{1h}$ into its original order, and form implicit strata and variance-strata as described section 1.0. To facilitate the theoretical development, it is desirable for each PSU to be in only one variance-stratum. However, this is not directly true, since some PSUs have a positive selection probability in two variance-strata. To eliminate this problem, PSUs with positive selection probability in two variance-strata will be converted into two new PSUs each completely contained within each of the respective variance-stratum.

Let PSU i have a positive selection probability in variance-stratum k and variance-stratum $k+1$. The selection probability of the new PSU, p_{ik+1}, from PSU i, associated with variance-stratum $k+1$, is $(C_{hi} - kI_h)/I_h$, where C_{hi} is the cumulative MOS for PSU i, defined in section 1.0 and I_h is the stratum sampling interval, (i.e., p_{ik+1} is the part of p_i that is contained in variance-stratum $k+1$). Then, $p_{ik} = p_i - p_{ik+1}$ (e.g., if p_i =0.3 and 0.2 of the probability is in variance-stratum $k+1$ then p_{ik} =0.1 and p_{ik+1} =0.2).

Some PSUs are now physically on the frame twice. Since the frame ordering of the PSUs is not changing, this modification does not change any of the possible samples or their selection probabilities. So, estimates and their respective variances have not changed.

We now make the "locally-random" assumption within the P_k's defined to be the variance-strata described above. (i.e., it is assumed that within each P_k, the original $\mathbf{F}_{1h}$ ordering represents one realization of the MOS randomization process described in section 2.3.)

2.5 A Discussion of ρ_{hij}

It has been assumed that there are multiple realization of a frame, either through a response variance model or super-population model. For any fixed frame realization, denoted by r, the correlation between x_i / p_i and x_j / p_j, ρ_{hij}^r has a value. Under the "locally-random" assumption, the stratum covariance, $C_h(x)$, equals $1/R\sum_{r=1}^{R}\sum_{i=1}^{n_h}\sum_{j\neq i}^{n_h}\rho_{hij}^r\sqrt{V_r(x_i / p_i)V_r(x_j / p_j)} = 1/R\sum_{r=1}^{R} C_h^r(x)$, where R is the number of frame realizations and assuming each frame realization are equally likely. This is just the average covariance taken across all possible randomizations. This averaging should reduce the number of extreme positive and extreme negative covariances and make the zero covariance assumption more acceptable.

2.6 Estimating $V(\hat{T}_{sy})$ and $v(\hat{T}_{sy})$

To estimate $V(\hat{T}_{sy})$, two assumption are made: 1) the "locally-random" assumption, described in 2.3, within the partitioning described in 2.4; and 2) $C_h(x)$ =0 or is small enough to be ignore.

With these assumptions, results from Kaufman (1999) provide an estimate for $V(\hat{T}_{sy})$ and $v(\hat{T}_{sy})$. Namely, $\hat{V}(\hat{T}_{sy}) = \sum_{h}^{H} \sum_{vs\in h}\left((\sum_{g\in vs} N_g^2 - N_{vs})/(N_{vs}(N_{vs}-1)) \right)\times \left(\sum_{i=1}^{N_{vs}} x_i^2 / p_i - T_{vs}^2 \right)$, where N_g, represents the number of PSUs in implicit stratum g with a positive selection probability, $N_{vs} = \sum_{g\in vs} N_g$ for variance-stratum vs, and $T_{vs} = \sum_{g\in vs} T_g = \sum_{g\in vs}\sum_{i=1}^{N_g} x_i$.

If i is a PSU in variance-strata k that has been changed into two PSUs because of positive selection probability in multiple variance-strata, it is assumed that $x_i = x_i' \times p_i / p_i^o$, where x_i' is the original data collected from PSU i and p_i is the new selection probability for the variance-stratum. As defined in section 2.4, p_i would be either p_{ik} or p_{ik+1}, depending on the variance-stratum. p_i^o is the original selection probability for i before the PSU was converted into two PSUs (i.e., $p_i^o = p_{ik} + p_{ik+1}$).

$$v(\hat{T}_{sy}) = \sum_{h}^{H} \sum_{vs\in h}\left(\left((\sum_{g\in vs} N_g^2 - N_{vs})/(N_{vs}^2 - \sum_{g=1}^{2} N_g^2) \right)\times \left(\sum_{i=1}^{2} 1/2(2x_i / p_i - \hat{T}_{vs})^2 \right)\right) \qquad (2)$$

where $\hat{T}_{vs} = \sum_{i\in vs} x_i / p_i$ is the sample estimate of T_{vs}.

If i is a PSU in variance-strata k that has been changed into two PSUs then $x_i / p_i = x_i' \ (p_i / p_i^o)\ (1/p_i) = x_i' / p_i^o$. Hence, $v(\hat{T}_{sy})$ can be estimated without knowing p_{ik} or p_{ik+1}.

The second term of the product in (2) is the balanced half-sample variance estimate (BHR) for the variance-stratum. Therefore, any differences between (2) and BHR can be attributed to the first term (scaling term).

The scaling term acts as an FPC. If the N_g's are all equal in a stratum then this term resembles the simple random sample FPC. However, when the stratum PSUs are skewed in either direction, this term can be greater that 1. In this situation, the BHR estimator should be expected to underestimate the variance.

When using the "locally-random" assumption, it is best to start with an initial ordering that will minimize the scaling factor. It is minimized when the number of PSUs in the respective implicit strata are equalized (i.e., equalizing the values of the MOS). Sorting the MOS in a serpentine manner can equalize the values of the MOS in a variance-stratum; thereby equalizing the number of PSUs within the implicit strata for a variance stratum.

2.7 Summary

With respect to the estimating the variance under the "locally-random" assumption, with this assumption along with assuming $C_h(x)$ =0, the above arguments provides an appropriate variance estimator for PPS systematic sampling, including an appropriate FPC. With respect to the estimating the variance under the fixed-ordered design variance-stratum assumption, as described in the introduction, even if the assumption is correct, the variances may still not be appropriate, because of an incorrect variance-stratum variance estimator.

3.0 Simulation

To compare the "locally-random" assumption with the fixed-ordered design variance-stratum assumption, a simulation study is presented comparing the BHR, under the fixed-order variance-stratum assumption; and the bootstrap variance estimator, under the "locally-random" assumption. The simulations will model the NCES's National Study of Postsecondary Faculty (NSOPF) Institution survey sample design. The frame will initially be ordered according to a variation of the NSOPF sample design. Under the fixed-ordered design assumption, the frame will not be randomized, while under the "locally-random" assumption, the frame will be randomly ordered, given the NSOPF frame ordering. $C_h^r(x)$ and $C_h(x)$ under the variance-stratum and "locally-random" assumption, respectively, will be realistic representations of the survey.

To be able to compute estimates for any selected sample, frame variables will be used in the simulation.

3.1 Sample Design

The first-stage sampling frame for NSOPF consisted of the 3,396 postsecondary institutions that were public or private not-for-profit Title IV participating institutions and provided formal degree programs of at least two years' duration. The 3,396 institutions in the NSOPF frame were stratified based on the highest degrees they offered and the amount of federal research dollars they received. Before sample selection, the frame was ordered by the previous NSOPF round's stratification, state and control number. A PPS systematic sample of 960 institutions was then selected using total faculty as the MOS. In order to fit this design into the framework of this paper, the frame ordering was changed. In this simulation study, the frame was ordered by the previous NSOPF round's stratification, state and MOS using a serpentine ordering as the state changes. The second stage sample of faculty is not used in this simulation, so it will not be described.

With almost 30% of the institutions in sample, this design provides a good check for an appropriate FPC.

3.2 Simulating the "Locally-random" Model

In practice, the original frame ordering is used for sample selection, but it is assumed that within variance-strata the original ordering of PSUs is one realization of the random model described in section 3.2. It is not necessary to specify the random ordering model. It is sufficient to know that there exists an underlying model. However, to do a simulation, it is necessary to specify the random ordering model, because the simulations need to use realizations from that model. The "locally-random" model is specified as a set, $\mathbf{R}$, which includes all possible random orderings, from which the original ordering is a member. $\mathbf{R}$ is defined as follows: First the frame is partitioned as specified in section 2.4, including the step which converts a PSU with positive selection probability in multiple variance-strata into two PSUs. This new frame is denoted by $\mathbf{F}_{2h}$. $\mathbf{R}$ is the set of all possible ordered frames that can be produced from all independent variance-stratum randomization of $\mathbf{F}_{2h}$. The first step in each simulation is to select a simulation sample from a randomly chosen $\mathbf{F}^{\circ}_{2h} \in \mathbf{R}$.

3.3 Bootstrap Variance (Locally-random Frame)

Kaufman (1999) provides a bootstrap variance estimator $V^*(\hat{T}_{sy})$ for (2) and it is used to estimate the locally-random assumption variance.

3.4 BHR Variances

The r^{th} school half-sample replicate is formed using the usual textbook methodology (Wolter, 1985) with 2 PSUs per stratum. When $n_h \geq 2$, PSUs are placed in variance-strata (see section 1.0), which are used as real strata for estimating variances. This is the BHR without FPC (BHR No FPC) variance. A second BHR variance estimate (BHR Prob FPC) 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. A third BHR variance estimate (BHR Random FPC) adjusts the first variance estimator by the variance-stratum scaling terms provide in (2). A fourth BHR variance estimate (BHR SRS FPC) adjusts the first variance estimator by $1 - n_h / N_h$. These BHR estimators are used to estimate the fixed-ordered frame variance.

3.5 Comparison Statistics

Below the statistics used to compare the "locally-random" assumption and the fixed-ordered design variance-stratum assumption variances are described.

3.5.1 Total Relative Covariance (TRC)

The total relative covariance for the estimated total of x, $\hat{T}_{sy}$ is estimated in the simulation by:

$$TRC = \left[\left(\sum_h \hat{C}_h(x) \right) / V(\hat{T}_{sy}) \right] \cdot 100 =$$

$$\left[\left(V(\hat{T}_{sy}) - \sum_h \sum_{g \in h} V(\hat{T}_{sy\,g}) \right) / V(\hat{T}_{sy}) \right] \cdot 100,$$

where: $V(\hat{T}_{sy}) = (1/(n_s - 1)) \sum_{s=1}^{n_s} (\hat{T}^s_{sy} - \overline{T}_{sy})^2$, $\hat{T}^s_{sy}$ is the estimated total from simulation sample s, and $\overline{T}_{sy}$ is the average of the $\hat{T}^s_{sy}$s over the n_s simulations. $V(\hat{T}_{sy\,g}) = (1/(n_s - 1)) \sum_{s=1}^{n_s} (\hat{T}^s_{sy\,g} - \overline{T}_{sy\,g})^2$, where: $\hat{T}^s_{sy\,g}$ is the sample estimated total for implicit stratum g, and $\overline{T}_{sy\,g}$ is the average of the $\hat{T}^s_{sy\,g}$ s over the n_s simulations. When sy =1 then s is selected from $\mathbf{F}_{1h}$ and when sy =2 then s is selected from $\mathbf{F}^{\circ s}_{2h}$, the randomly selected frame for s.

3.5.2 Relative Error including TRC

Rel. Error $= \left[\left(\overline{V}^e(\hat{T}_{sy}) + TRC \right) / V(\hat{T}_{sy}) - 1 \right] \cdot 100$, where $\overline{V}^e(\hat{T}_{sy})$ is the average of the variance estimates ($V^e(\hat{T}^s_{sy})$) across the simulation samples denoted by s. When sy =1, the BHR procedure ($e = 1$) is used and when sy =2 the bootstrap procedure ($e = 2$) is used.

For averages, the relative error and total relative covariance are computed as described above, by replacing $\hat{T}^s_{sy}$ and $\hat{T}^s_{sy\,g}$ with the respective averages.

3.6 Number of Replicates and Simulations

Forty-four and forty-five replicates have been used in the BHR and bootstrap variances, respectively. Total relative covariances are computed using 3,000 simulations and Bootstrap and BHR variances are computed using 500 simulations.

4.0 Results

Table 1 provides the percent distribution of total relative covariance for the eighteen estimates used in the simulation. It is clearly incorrect to assume the covariances are zero. Most of the time, the covariances are negative, which indicates that the variances should be overestimates. For the fixed-ordered design (row 1), 16.7% of the time the covariances are smaller than –30% with the smallest being –63%. In many situations, these would represent an unacceptably large overestimation. Since each covariance produced from the randomly-ordered design (row 2) is an average of all possible fixed-ordered design covariances, one would expect fewer extreme relative covariances. This is true with no covariance less than -30%.

Some of the covariances are positive. From row 1, 16.6% of the covariances are positive. The largest is 56%, which in most situation, will produce too large of a variance underestimation. From row 2, 5.6% of the covariances are positive. Again, as expected, the magnitude of the extreme covariances is reduced with the randomly-ordered design. The largest total relative covariance being 24%.

To determine the best FPC, it is necessary to include the total covariance term into the relative error estimate, because the total covariance term clearly can not be closely approximated by zero. Table 2 provides the percent distribution of the relative errors for the five variance estimators including the total covariance term.

The randomly-ordered design variances accurately estimate the true variance. Since there are no extreme estimates, the FPC is clearly appropriate.

For the fixed-ordered design, the No FPC adjusted variances overestimate the variance 72.2% of the time. It's not surprising that the overestimation can be large when 30% of the frame is in sample. What is surprising is that some of the variances are still underestimated, 27.8% of the time, even with the total covariance included in the relative error. Table 3 provides each of these individual estimates. The largest underestimate is almost –15%. This would indicate that the variance estimator based on a fixed-ordered design has some additional sources of error, such as the exclusion of an appropriate FPC that can be greater than one.

Including an FPC that can be greater than one, like the Random FPC, is empirically justified in table 3. By comparing the Random FPC and No FPC adjusted variances, it can be seen that the Random FPC adjusted variances reduce the underestimation in all but one case, while the SRS and Prob FPC adjusted variances always make the underestimation a great deal worst. The Random FPC adjusted variances produce a large increase in the absolute error only once, while the SRS and Prob FPC adjusted variances are always worst.

For the fixed-ordered design, table 2 shows that the Prob FPC adjusted variances underestimate the variance almost 78% of the time and produce a large underestimate almost 56% of the time. This indicates that the Prob FPC adjustment tends to produce too large of a variance reduction.

For a fixed-ordered design, the Random FPC adjusted variances performs well. There are no large underestimates and large overestimates are produced almost 28% of the time. This indicates that the random FPC seems to be reasonable, even though some of the FPCs are greater than 1.

For the fixed-ordered design, the SRS FPC adjusted variances produce a large underestimate 33% of the time and produce a large overestimate 5.6% of the time.

Of the four FPC adjusted variances using a fixed-ordered design, the Random FPC adjusted variance seems to perform the best. This may indicate, when selecting PPS systematic samples, that an FPC greater than 1 can be reasonable, even though the frame is not randomly-ordered.

5.0 Conclusions

There are a number of conclusions that can be drawn for these results. The first is that it is not necessarily wise to assume the total covariances are zero. Assuming they are zero can produce unacceptably large overestimates, as well as unacceptably large underestimates. However, it would be difficult to eliminate this assumption altogether without allowing for the possibility of negative variances. What seems reasonable is to design surveys to eliminate or reduce the number of extreme covariances. When the "locally-random" assumption is appropriate, the number of extreme covariances is reduced. Section 6.0 provides sample designs suggestions that can make the "locally-random" assumption more appropriate, thereby reducing the expected number of extreme covariances.

A second conclusion is when the "locally-random" assumption is appropriate and the total covariance is included into the relative errors, variances are correctly estimated with the randomly-ordered variance estimator. So, the scaling factor in (2) produces an appropriate FPC for the PPS systematic selection process.

The third conclusion is that when the "locally-random" assumption is not appropriate, the four FPCs used in this study produce mixed results. None of the FPCs evaluated work 100% of the time. This is not surprising because there is no theoretical justification for any of them. However, the Random FPC seems to work best for this particular sample design.

A fourth conclusion is when: 1) the "locally-random" assumption is not appropriate, 2) no FPC adjustments are made to the variance estimator, and 3) covariances have been added back in, variances can still be reasonably sized underestimates. This could indicate that an appropriate FPC may be greater than 1 at times.

6.0 Sample Design Implications

To improve variance estimation based on these results, a number of survey design modifications can be suggested. The main objective is to make the ordered frame look more "locally-random", while reducing the scaling factor in (2). These objectives can be accomplished by making the frame look more continuous with respect to the MOS, as one looks at the ordered frame from top to bottom within each stratum. If this can be accomplished, the MOS will be roughly equal within variance strata, which will make the "locally random" assumption more appropriate, and reduce the scaling factor. To do this: 1) make the MOS the last variable in the frame ordering process; 2) order the MOS in a serpentine manner, as other ordering variables switch values (This may imply categorizing continuous variables before ordering the frame.); and 3) reduce the number of ordering variables. All this can be done without physically randomizing the frame.

With respect to the fixed-ordered sample design used in this simulation, the above discussion suggests that state be eliminated from the frame ordering or maybe replaced by some sort of regional groupings. This was not tested in a simulation. However, in a preliminary simulation analysis, the unmodified ordering described in section 3.3 (i.e., no serpentine ordering of the MOS), was used. In those preliminary simulations, not presented here, the number of extreme total covariance terms increased by 74% compared to the fixed-ordered sample design simulation results presented here. So the serpentine ordering of the MOS is helpful.

An additional option is to physically impose the "locally-random" assumption by first randomly selecting an $\mathbf{F}^{\circ}_{2h} \in \mathbf{R}$ as described in section 3.1 and then selecting the sample from $\mathbf{F}^{\circ}_{2h}$. In this situation, the frame is "locally-random", so formula (2) provides the correct variance, assuming the total covariance is zero. Even if the total covariance is not zero, the number of extreme relative errors will be reduced. Since the variance strata are formed using the $\mathbf{F}_{1h}$ ordering, much of the original efficiency is maintained. The drawback is that it becomes possible to select some PSUs multiple times. It may then be necessary to increase the sample size in the strata to yield the original expected sample size; or one could randomize only PSUs that are selectable from a single variance stratum and avoid the multiple selection issue.

7.0 References

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

Kaufman, Steven (1999). "Using the Bootstrap to Estimate the Variance from a Single Systematic PPS Sample," *Proceedings for the Section on Survey Methods, American Statistical Association*, pp. 683-688. Alexandria, Va.: American Statistical Association.

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

Table 1 -- %Distribution of the relative total covariance

Design	Relative covariance $\le -30\%$	-30%<Relative covariance<0%	$0\% \le$ Relative covariance<15%	Relative covariance $\ge 15\%$	Minimum	Maximum
Fixed-ordered	16.7	66.7	11.0	5.6	-63	56
Randomly-ordered	0.0	94.4	0.0	5.6	-27	24

Table 2 -- % Distribution of relative error including the total covariance term

Design	Relative error $\le -15\%$	-15% <Relative error<0%	$0\% \le$ Relative error<30%	Relative error $\ge 30\%$	Minimum	Maximum
Randomly-ordered	0	22.2	77.8	0	-11	14
Fixed-ordered						
No FPC	5.6	22.2	16.6	55.6	-15	84
Random FPC	0	27.8	44.4	27.8	-14	49
SRS FPC	33.3	16.6	44.4	5.6	-34	30
Prob FPC	55.6	22.2	22.2	0.0	-46	6

Table 3 -- Relative error (%) including total covariance term

Variables	Randomly-ordered design (bootstrap)	Fixed-ordered design (Balanced half-sample replication)			
		No FPC	Random FPC	SRS FPC	Prob FPC
Number of Faculty on 11/12 Month Schedule	3.6	-7.3	-2.8	-21.0	-31.1
Salary Outlay of Faculty on 11/12 month Schedule	5.5	-9.7	-7.2	-24.9	-36.8
Salary Outlay of Faculty on 9/10 month Schedule	6.9	-5.1	-5.0	-21.8	-32.9
Number of Full-time Faculty	14.1	-14.5	18.4	-33.8	-46.1
Number of Men Faculty	5.0	-2.3	-14.0	-29.1	-45.1