Biometrika (1986), 73, 2, *pp.* 485-91
Printed in Great Britain

Some asymptotic results for the systematic and stratified sampling of a finite population

BY PHILLIP S. KOTT
Energy Information Administration, Washington, D.C. 20583, U.S.A.

SUMMARY

A sample mean has an asymptotically insignificant model bias under a broad set of functional relationships between the parameter of interest and an auxiliary size variable given a systematic sample from a size ordered list. When based on randomly drawn samples within size strata, this estimator of the population mean not only has an asymptotically insignificant model bias but is design consistent as well. In addition, there are more efficient estimation strategies that have insignificant model biases and may be design consistent.

Some key words: Auxiliary variable; Balanced sample; Design consistency.

1. INTRODUCTION

Suppose we believe that a parameter of interest, Y, for a finite population of N units is related to a known auxiliary variable, X, but we are unsure of the functional relationship between them. How should we go about estimating the population mean, $\bar{Y}_N$, based on a sample of n units?

There is extensive literature on this subject when the functional relationship between Y and X is assumed to be linear with unknown parameters. An excellent text with an extensive bibliography is Cassel, Särndal & Wretman (1977). Unfortunately, the exact relationship between the variables is often unknown and likely to be complicated.

Statisticians have imposed a variety of restrictions on the sampling design to protect themselves against the failure of the simple linear model. Among them are design unbiasedness (Godambe, 1955), asymptotic design unbiasedness (Brewer, 1979), and design consistency in the asymptotic sense of Isaki & Fuller (1982).

A fully model dependent approach is taken by Royall & Herson (1973). They argue that a sample mean, $\bar{Y}_S = \Sigma_S Y_i / n$, will be a robust estimator of a population mean if the sample is balanced not only on X, but also on powers of X up to degree J. These constraints on the sample remove the model bias from $\bar{Y}_S$ if Y_i is a Jth degree polynomial function of X_i plus a random error. What happens, however, if the real relationship between Y and X is not of this form? Suppose, for example, $E(Y_i) = \exp(X_i)$.

We will see that, when the sample is drawn systematically from a list ordered by X values, the sample is asymptotically balanced on any bounded monotonic transformation of X. In other words, when $E(Y_i)$ is a finite combination of bounded monotonic transformations of X_i, the model bias of the sample mean is an asymptotically insignificant contributor to model mean squared error. This broad theoretical result has, of course, practical limitations to be discussed.

No estimation strategy involving systematic sampling is design consistent in the Isaki-Fuller sense. By replacing the systematic sample above with an analogous stratified sample, we can retain the desirable model based property when Y depends on X alone, while assuring design consistency to protect against the possibility of bias due to the existence of additional variables upon which Y may also be functionally dependent.

The problem with using the sample mean as an estimator is that it is not very efficient. Fortunately, the asymptotic properties discussed above can be incorporated into more efficient estimation strategies by employing probability proportional to size sampling designs or a model-based stratification rule.

2. The basic model

Let P be a population of N members. For each member, U_i $(i=1,\ldots,N)$, let Y_i be specified by the model

$$Y_i = \beta X_i + \varepsilon_i V_i, \tag{2·1}$$

where the ε_i are independent and identically distributed with mean zero. The X_i in (2·1) are known, while the V_i need not be.

We are interested in the properties of $\bar{Y}_S = \Sigma_S Y_i/n$, where S denotes a sample of n units, as an estimator of $\bar{Y}_N = \Sigma_P Y_i/N$. The model mean squared error of $\bar{Y}_S$ is

$$E_\varepsilon\{(\bar{Y}_S - \bar{Y}_N)^2\} = \beta^2(\bar{X}_S - \bar{X}_N)^2 + \sigma_\varepsilon^2\left(\frac{1}{n^2} - \frac{2}{nN}\right)\sum_S V_i^2 + \sigma_\varepsilon^2 N^{-2}\sum_P V_i^2. \tag{2·2}$$

The first term in (2·2) is the contribution of the model bias of $\bar{Y}_S$ to its model mean squared error. The remaining terms, the model variance of $\bar{Y}_S$, is the contribution of model standard deviation to model mean squared error.

We will allow both the population size, N, and the sample size, n, to grow arbitrarily large with $k = N/n$ a fixed integer greater than one. To put this in a more formal setting, consider an infinite sequence of populations, $P_1 \subseteq P_2 \subseteq \ldots$, with $N_j = kn_j$ units in population P_j, and suppose that the sequence $\{N_j\}$ is unbounded. If $n = n_j$ and $N = kn$, then n and N becoming arbitrarily large is equivalent to j becoming arbitrarily large.

We use the following boundary conditions in our subsequent analysis:

$$|X_i| < B_x < \infty, \quad 0 < B_1 < V_i < B_2 < \infty. \tag{2·3}$$

It is critical that X_i not be allowed to grow arbitrarily large with n.

The model standard deviation of $\bar{Y}_S$ is bounded below by $\sigma_\varepsilon B_1 \surd\{(1-1/k)/n\}$ and above by $\sigma_\varepsilon B_2 \surd\{(1-1/k)/n\}$. Thus its contribution to model mean squared error is of order n^{-1}. The order of the contribution of the model bias depends on the sample itself. If the sample is balanced on X, that is $\bar{X}_S = \bar{X}_N$, then the model bias offers no contribution to the model mean squared error.

On the other hand, suppose the sample is chosen via simple random sampling without replacement, and $\Sigma(X_i - \bar{X}_N)^2/N$ converges to a finite positive value. Consider the combined mean squared error of $\bar{Y}_S$, $E_\varepsilon[E_D\{(\bar{Y}_S - \bar{Y}_N)^2\}]$, where E_D is the expectation relative to the sampling design. Applying this expectation operator to both sides of (2·2), we see that the contribution of the model bias of $\bar{Y}_S$ to its combined mean squared error is $\beta^2 \operatorname{var}_D(\bar{X}_S)$, which is of order n^{-1}. The contribution of the model standard deviation is of the same order.

3. Size ordered stratification

There is a middle ground between forcing samples to be balanced on X and choosing the sample randomly. Let the population be ordered by X values and then divided into n strata of equal size. In other words, relabel the units so that $X_1 \leqslant \ldots \leqslant X_N$, and let the set $\{U_1, \ldots, U_k\}$ form stratum 1, $\{U_{k+1}, \ldots, U_{2k}\}$ form stratum 2, and so forth. After stratifying the population, let one unit from each stratum be chosen for the sample. At this point, we are setting no restrictions on how the units are chosen within strata aside from requiring that the process not depend directly on the size of the Y_i.

Note that the X_i are relabelled for each population in the sequence $\{P_j\}$. As a result, X_i in P_j need not equal X_i in P_k. Moreover, suppose S_j denotes the sample from population j. The sequence of samples, $\{S_j\}$, will not in general be nested; that is $S_1 \subseteq S_2 \subseteq \ldots$ need not hold.

Given the boundary condition in (2·3) and the sampling design just described, $\bar{X}_S$ at its minimum is $(X_1 + X_{k+1} + \ldots + X_{(n-1)k+1})/n$. At its maximum, it is

$$(X_k + X_{2k} + \ldots + X_N)/n.$$

Since $X_{mk} \leq X_{mk+1}$ for all m, the difference between $\bar{X}_S$ at its minimum and its maximum is bounded by $(X_N - X_1)/n$ and $2B_X/n$. When $\bar{X}_S$ is averaged over all possible samples the result is $\bar{X}_N$. Therefore $(\bar{X}_S - \bar{X}_N)$ must be bounded by $2B_X/n$. As a result, if the above size ordered stratified sampling plan is followed, the contribution of the model bias of $\bar{Y}_S$ to its mean squared error, as expressed in (2·2), is bounded by $4\beta^2 B_X^2/n^2$. As N grows arbitrarily large, this contribution tends to zero faster than the model mean squared error itself, which must be of order n^{-1} due to the contribution of the model standard deviation.

We have not yet specified a mechanism for choosing the sampled units from each stratum. One popular method, systematic sampling, involves selecting an integer, t, between 1 and k and then choosing $U_t, U_{t+k}, U_{t+2k}, \ldots, U_{t+(n-1)k}$ for the sample. Exactly one unit from each stratum is sampled using this approach. The start point, t, may or may not be chosen randomly.

Another approach is to sample one unit from each stratum randomly, using equal selection probabilities within strata. Both this sampling design and systematic sampling with a random start point results in a design unbiased estimator; that is $E_D(\bar{Y}_S - \bar{Y}_N) = 0$. Design consistency is another matter, as we see later.

4. Monotonic transformations

We call any sample from a sampling design using the kind of size ordered stratification described last section asymptotically balanced on X if the $|X_i|$ are bounded. This terminology derives from the fact that $\bar{X}_S - \bar{X}_N$ approached zero faster than the model standard deviation of $\bar{Y}_S$. Not only is such a sample asymptotically balanced on X, it is also asymptotically balanced on any bounded monotonic transformation of X. In fact, it is asymptotically balanced on any finite combination of bounded monotonic transformations of X.

Before proving this assertion, a word of caution is in order about its apparent generality. Since Y and X take on N discrete values, it is possible to map Y onto X with a continuous function containing no more than $N-1$ turning points. Thus Y itself is a finite combination of bounded monotonic transformations of X. In the context of sampling from a finite population, however, the asymptotic property asserted above is relevant only when the number of transformations is much smaller than the sample size, not just the population size. Consequently, if it is suspected that there are more than one or two turning points in the graph of $E_\varepsilon(Y)$ against X consideration of another auxiliary variable may be wise.

For the proof, let Y_i be specified as

$$Y_i = \sum_{j=1}^{J} \beta_j Z_{ji} + \varepsilon_i V_i, \tag{4·1}$$

where the ε_i and V_i are as before, $Z_{ji} \geq Z_{jm}$ if $X_i \geq X_m$, and $|Z_{ji}| < B_Z$ for all j. Since each Z_j is a monotonic transformation of X, the sampling design is the same as if Z_j were the auxiliary variable not X. Because each Z_j is bounded, all $\bar{Z}_{js} - \bar{Z}_{jN}$ are of order n^{-1}. There are only a finite number of $\bar{Z}_{js} - \bar{Z}_{jN}$; consequently, the model bias of $\bar{Y}_S$ under the model in (4·1), $\Sigma\beta_j(\bar{Z}_{jS} - \bar{Z}_{jN})$, must be of order n^{-1}. This model bias is an asymptotically insignificant contributor to the model mean squared error, the latter being dominated by model standard deviation, which has not changed.

The use of models to analyse systematic and stratified sampling schemes dates back at least to Cochran (1946). If the problem were simply to limit the bias under the model, $Y_i = \beta X_i + \varepsilon_i$, then using the ratio estimator, $\hat{Y}_{RAT} = \bar{X}_N \bar{Y}_S / \bar{X}_S$, with any sample would provide an obvious model unbiased solution.

We have seen that systematic and stratified sampling offer more robust bias protection than the ratio estimator per se. As it happens, the ratio estimator when based on such samples possesses the asymptotic model properties of $\bar{Y}_S$, because $(\bar{Y}_S - \hat{Y}_{RAT})$ is of order n^{-1} when $|\bar{X}_S|$ is bounded from below by a positive number.

5. Design consistency

Suppose the population were such that the Y_i could be expressed as

$$Y_i = X_i + (-1)^i, \tag{5·1}$$

where $X_i = 1 - 2^{-i}$. Consider the asymptotic properties of $\bar{Y}_S$ based on a systematic sample of units drawn from this already X ordered population.

Holding $k = N/n$ constant at an integer value, observe that as n approaches infinity: (i) $\lim \bar{Y}_N = 1$; (ii) $\lim \bar{Y}_S = 1$ if k is odd; (iii) $\lim \bar{Y}_S = 0$ if k is even and t, the start point, is odd; (iv) $\lim \bar{Y}_S = 2$ if both k and t are even. The important thing is that when k is even, $\bar{Y}_S - \bar{Y}_N$ does not approach zero as N grows arbitrarily large.

It is generally impossible to construct a sampling design, D, so that $\bar{Y}_S - \bar{Y}_N$ converges to zero asymptotically no matter how the Y_i are distributed in the population. It is often possible, however, to construct a D so that the probability of choosing a sample for which $\lim(\bar{Y}_S - \bar{Y}_N)$ does not tend toward zero is arbitrarily small.

We say that an estimation strategy for $\bar{Y}_N$, a couple $\hat{Y}$ and D, are design consistent if the design mean squared error of $\hat{Y}$ tends to zero when k is any finite integer and n grows arbitrarily large; mathematically, $\lim E_D\{(\hat{Y} - \bar{Y}_N)^2\} = 0$. Systematic sampling, even with a random start point, is not design consistent, as we can see in the example above.

On the other hand, suppose the population is stratified as before, and one unit is sampled randomly from each stratum using equal selection probabilities within strata. Call this design D_{STR}. The design mean squared error of $\bar{Y}_S$ is $\Sigma_u \Sigma_v (Y_{uv} - \bar{Y}_{u.})^2/(kn^2)$, where $Y_{uv} = Y_{v+(u-1)k}$, and $\bar{Y}_{u.} = \Sigma_v Y_{uv}/k$. If the $|Y_i|$ are bounded, then the design mean squared error of $\bar{Y}_S$ is of order n^{-1}, and $(\bar{Y}_S, D_{STR})$ is design consistent.

In terms of combined mean squared error, the design consistency of $(\bar{Y}_S, D_{STR})$ assures that even if Y is functionally dependent on another auxiliary variable in addition to X, the combined mean squared error of $\bar{Y}_S$ becomes smaller as the sample size grows larger.

The definition of design consistency used here is mathematically equivalent to that of Isaki & Fuller (1982). It is very different from that of Cochran (1977, p. 21), which simply lets n tend toward N. Any estimation strategy requiring the sampled units to be distinct coupled with an estimator of the form $\bar{Y}_S + G(\bar{X}_S - \bar{X}_N)$, where G is a bounded function of the sample, is design consistent in Cochran's sense, no matter how unlikely the G. Moreover, while the sample size never becomes arbitrarily large in practice, often n is large enough for asymptotic results to have practical consequences. On the other hand, the sample size never approaches the population size in practice, except in a complete census.

Another meaning of design consistency is attributed to Brewer by Little (1983), who uses the term 'asymptotic design consistency'. Brewer (1979) suggests replicating a finite population r times and selecting independent samples in each replicate using the sample design. As r grows arbitrarily large, the design bias and mean squared error of the estimator in question can be analysed as an estimator for the pooled mean, $\Sigma\Sigma Y_{ki}/(rN)$, where Y_{ki} is the ith population unit of replicate k.

Because the sample is chosen independently in each replicate in Brewer's approach, it fails to capture the properties of size ordered stratification for large samples in a meaningful way. In the Isaki-Fuller approach followed here, replicating the population any number of times may be the formal mechanism generating arbitrarily large sample and population sizes, although one is not confined to this mechanism. The units sampled in each replicate, however, may change as the number of replicates grows, because rather than drawing an independent sample from each replicate, one draws a single sample from the pooled population.

6. More efficient estimation

So far we have concerned ourselves with model bias and design consistency but not with model or design efficiency. In this section, we will extend the analysis of the previous sections to estimation strategies with smaller model variances.

Recall that the model standard deviation of Y_i in (2·1) is V_i. If $V_i = X_i > 0$ for all i, and no V_i exceeds $N\bar{V}_N/n$, then the linear design unbiased estimation strategy with the least combined mean squared error is the Horvitz-Thompson estimator, $\hat{Y}_{\text{HT}} = \bar{V}_N\bar{Q}_S$, where $Q_i = Y_i/V_i$, coupled with probability proportional to V_i sampling (Godambe, 1955). Moreover, if again $V_i = X_i > 0$ for all i, then, under the model

$$Q_i = \beta_1 + \beta_2 V_i + \varepsilon_i, \tag{6·1}$$

the most model efficient of all model unbiased estimation strategies combines $\hat{Y}_{\text{HT}}$ with a sampling design restricted to π-balanced samples; i.e. samples satisfying $\bar{V}_S = \Sigma_P V_i^2/\Sigma_P V_i$ (Kott, 1984).

If $nV_i/(N\bar{V}_N)$ is bounded by unity, we can show that a systematic probability proportional to V_i sample (Madow, 1949) from a V ordered list is asymptotically π-balanced; under the model in (6·1), the model bias of $\hat{Y}_{\text{HT}}$ given this sample is asymptotically insignificant as a contributor to model mean squared error.

To draw such a sample, we rearrange the population so that $V_1 \leqslant \ldots \leqslant V_N$, and calculate for each i cumulative V sums, $CV_i = \Sigma V_j$, where j runs from 1 to i. Then n hit points, $H_1, \ldots, H_n$, are determined by choosing a start point, t, on the interval (0, 1] and letting $H_k = \{t + (k-1)\}CV_N/n$. The sample consists of those units i such that $CV_i \geqslant H_k > CV_{i-1}$.

Any start point, whether it is randomly chosen or not, will generate a sample such that $\hat{Y}_{\text{HT}}$ has an insignificant model bias under (6·1) when V is bounded. This is because the model bias of $\hat{Y}_{HT}$ is bounded by $\bar{V}_N\beta_2(V_N - V_1)/n$ and thus of order n^{-1}; while the model standard deviation of $\hat{Y}_{HT}$ is of order $n^{-\frac{1}{2}}$.

In fact, if $E_\varepsilon(Q_i)$ is any finite combination of bounded monotonic transformations of V_i, and $\text{var}_\varepsilon\,(Q_i) = \sigma_\varepsilon^2$, the Horvitz-Thompson estimator coupled with a systematic probability proportionate to V_i sample from a V ordered list has an asymptotically insignificant model bias and an asymptotic model variance equal to the minimum attainable under (6·1).

In practice, the V_i are rarely known. It is often assumed that the V_i have the form X_i^γ, where γ is between 0 and 1, given the X_i are all positive. Choosing a value for a particular application is a dubious exercise even when based on previous experience. Nevertheless, it is almost always preferable, in terms of minimizing model variance, to assuming $\gamma = 0$, which is what using $\bar{Y}_S$ based a systematic sample does implicitly. Cochran (1977, p. 257) states the commonly held belief that $\frac{1}{2} < \gamma < 1$ for most applications.

When the choice for each V_i does take the form X_i^γ, $0 < \gamma \leqslant 1$, all finite combinations of bounded monotonic transformations of X can be expressed as finite combinations of bounded monotonic transformations of V. Moreover, $E_\varepsilon(Q_i) = X_i^{1-\gamma} E_\varepsilon(Y_i/X_i)$. As a result, if $E_\varepsilon(Y_i/X_i)$ can be expressed as a finite combination of bounded monotonic transformations of X_i, so can $E_\varepsilon(Q_i)$. This allows us to claim the following: If $V_i = X_i^\gamma$, and $E_\varepsilon(Y_i/X_i)$ is a finite combination of bounded monotonic transformations of X_i, then $\hat{Y}_{HT}$ based on a systematic probability proportional to V_i sample from a V ordered list has an insignificant model bias. This is true whether or not $\text{var}_\varepsilon\,(Y_i) = V_i^2\sigma_\varepsilon^2$.

Given a systematic probability proportional to V_i sampling design from a V ordered list, D_{PSY}, using a start point chosen randomly from a uniform distribution on (0, 1], $\hat{Y}_{HT}$ is design unbiased. The estimation strategy $(\hat{Y}_{HT}, D_{PSY})$ is not design consistent however. Equation (5·1) with $Y_i = Q_i$ and $X_i = V_i$ serves as a counterexample. One may calculate lim $\hat{Y}_{HT}$ when $(N-1)/n$ is held at an even value and compare it to lim $\bar{Y}_N$.

If it were possible, we would divide the V ordered population into n strata with equal cumulative V sizes; that is $\Sigma_{(1)} V_i = \ldots = \Sigma_{(n)} V_i$ with (u) denoting the elements of stratum u, and then randomly select one unit for each stratum using probability proportional to to V_i sampling. This design, D_{PST}, combines with $\hat{Y}_{HT}$ to be design consistent when the Q_i are bounded, while retaining the model properties of $(\hat{Y}_{HT}, D_{PSY})$.

It is often impossible to divide the population into equally sized strata. As a result, some model efficiency has to be sacrificed to attain design consistency. Let W_u be the share of V in stratum u, $W_u = \Sigma_{(u)} V_i/(N\bar{V}_N)$, where all the W_u are as nearly equal as possible. Let $\hat{Y}_{SH}$ be the stratified Horvitz-Thompson estimator, $\bar{V}_N \Sigma_u W_u \bar{Q}_S$, where $\bar{Q}_S$ equals the Q_i value for the sampled unit in stratum u. Let D_{PST} again be the sampling design: one unit to be chosen from each stratum by probability proportional to V_i sampling. The estimation strategy $(\hat{Y}_{SH}, \hat{D}_{PST})$ is design consistent if the Q_i and the nW_u are bounded as n becomes arbitrarily large. In addition, when $E_\varepsilon(Q_i)$ is a finite combination of bounded monotonic transformations of V_i, the model bias of $\hat{Y}_{SH}$ is asymptotically insignificant.

7. Some remarks on stratified sampling

The focus thus far has been restricted to stratified samples with one sampled unit per stratum. It is simple to show that the asymptotic properties of $\bar{Y}_S$ and $\hat{Y}_{SH}$ remain when an equal, bounded number of units are sampled per stratum. Equality of stratum sample sizes is also unnecessary when the estimators are adjusted in an appropriate manner. An example of this is $\hat{Y}_{SS} = \Sigma(N_u/n_u)\Sigma Y_i/N$, where the second sum runs over the randomly selected units of u. Even the N_u in $\hat{Y}_{SS}$ may vary as long as the population size of each stratum is bounded.

A word of caution is in order when using any of these extensions: one must be careful that there are enough strata, all of sufficiently small size, to justify the invocation of asymptotic properties.

It is often required to estimate the mean squared errors of estimators. Since the V_i in (4·1) are in general unknown, estimating the model mean squared error of an estimator of $\bar{Y}_N$ derived using this model would require some heroic, and perhaps erroneous, assumption.

An alternative strategy in this situation is to use $\hat{Y}_{SS}$ based on two randomly selected units per stratum and estimate the design variance of $\hat{Y}_{SS}$ in the usual way. If one believes, a priori, that $V_i = F(X_i)$ for some monotonic F, then the population can be stratified so that the $N_u \Sigma F(X_i)$, where the sum runs over the population units in u, are roughly equal. Under this model-based stratification scheme, which approximates the probability proportional to V_i sampling of the last section, the design expected model variance of $\hat{Y}_{SS}$ is minimized when the stratum sample sizes are equal. By setting all the n_u equal to 2, we achieve the largest number of strata that allows design unbiased variance estimation. Thus one's preconceived notion of the V_i can be used in a robust manner in the design of the population mean estimator itself without affecting the estimation of its design variance.

The model bias of $\hat{Y}_{SS}$ is small, but with a finite sample not ignorable. As a consequence, a slightly better estimator than $\hat{Y}_{SS}$, one that also has a well known and indisputable design variance estimator, may be the difference estimator, $\hat{Y}_D = \hat{Y}_{SS} - B(\hat{X}_{SS} - \bar{X}_N)$, where B is an a priori notion of the average value of $dE_\varepsilon(Y)/dX$. By estimating B from the sample, an even better estimator may be constructed. Unfortunately, the variance estimator of this generalized difference estimator, while design consistent, has a relative design bias of order n^{-1}. If one is unhappy with $\hat{Y}_{SS}$ because of its model bias of order n^{-1}, one may be unequally unhappy with the variance estimator of the generalized difference estimator.

References

Brewer, K. R. (1979). A class of robust sampling designs for large-scale surveys. *J. Am. Statist. Assoc.* **74**, 911-5.

Cassel, C. M., Särndal C. E. & Wretman, J. H. (1977). *Foundations of Inference in Survey Sampling.* New York: Wiley.

Cochran, W. G. (1946). Relative accuracy of systematic and stratified random samples for a certain class of populations. *Ann. Math. Statist.* **17**, 164-77.

Cochran, W. G. (1977). *Sampling Techniques*, 3rd ed. New York: Wiley.

Godambe, V. P. (1955). A unified theory of sampling from finite populations. *J. R. Statist. Soc.* B **17**, 269-78.

Isaki, C. T. & Fuller, W. A. (1982). Survey design under the regression superpopulation model. *J. Am. Statist. Assoc.* **77**, 89-96.

Kott, P. S. (1984). A fresh look at bias-robust estimation in a finite population. *Proc. Survey Res. Meth. Sect., Am. Statist. Assoc.*, pp. 176-8.

Little, R. J. (1983). Estimating a finite population mean from unequal probability samples. *J. Am. Statist. Assoc.* **78**, 596-604.

Madow, W. G. (1949). On the theory of systematic sampling, II. *Ann. Math. Statist.* **20**, 333-54.

Royall, R. M. & Herson, J. (1973). Robust estimation in finite populations, I. *J. Am. Statist. Assoc.* **68**, 880-9.

[*Received October* 1984. *Revised November* 1985]