Computing Variances from Complex Samples with Replicate Weights

Cathryn S. Dippo, Bureau of Labor Statistics
Robert E. Fay, Bureau of the Census
David H. Morganstein, WESTAT Inc.

1. INTRODUCTION

The literature on the estimation of design-based sampling variances for complex sample designs emphasizes basically two principal strategies for non-linear estimators: linearization and replication methods. Several considerations enter into the choice between these two methods; three of these have individually counted against the choice of replication:

i) Many sample designs do not satisfy the restrictions required by the most familiar replication techniques, while linearization provides a general approach whenever an estimator of variance for estimates of population total exists and the statistics of interest are smooth functions of such totals.

ii) In application, more precision in the variance estimate, through consideration of more "degrees of freedom," is sometimes achieved with linearization.

iii) Replication methods may require more extensive computation.

Recent theoretical and practical developments mitigate or offset these considerations, increasing the attractiveness of replication:

i) Although balanced repeated half-sample replication, jackknifing, and random groups cover a limited set of situations in their original form, the literature has evolved a large number of adaptations of these methods to more specific situations, such as the reflection of finite population corrections and multiple stages of selection. The result of Fay (1984) simply states the logical conclusion of this line of development: there is no variance estimator based on sums of squares and cross-products that cannot be represented by a resampling plan.

ii) Given the potential use of general resampling plans (e.g. Fay 1984), the reliability of a variance estimate from replication for a linear estimator relative to its linearization counterpart is limited only by the number of replicate computations employed. Since it is feasible to do 50, 100, or even 200 such computations, the practical implications of this second criticism are relatively unimportant in many cases.

iii) It is true that replication methods will continue to be "computer-intensive" (Diaconis and Efron 1983). In many applications, however, implementation of replication methods may be far less "person-intensive" than methods based on linearization, in the sense that replication may require far less professional time to implement in new situations. Furthermore, replication may facilitate the estimation of variance for highly complex functions for which linearization becomes a practical impossibility. The implementation of replication through replicate weights may facilitate the computation of design-based variances from a given data set by far more researchers for their own applications.

Our paper is essentially a detailed discussion of the last of these three points. Specifically, we will discuss the simplification in implementing replication methods by assigning each survey case a series of replicate weights. This is not particularly a new idea: this paper will describe an application of this method by the first author to a survey in 1973, and the basic notion has possibly been independently discovered by other researchers. The objective of this paper is to review this technique and to emphasize the favorable practical experience that each of the authors has acquired with this approach.

2. BASIC APPROACH

In almost all surveys, weights assigned to each sample case play a key role in estimation. Typically, a single final weight W_i is computed for each survey case i. (Some surveys have multiple sets of weights for estimation, and generally the approach described here can be adapted to this situation as well.) For any characteristic X_i defined for each survey case i, an estimate X_0 of the population total for this characteristic is given as the weighted sum of the characteristic over the survey cases

$$X_0 = \sum_i W_i X_i \tag{2.1}$$

Estimates of proportions, means, ratios, regression coefficients, etc. are almost always functions of estimates in the form of (2.1) (including, for example, the complex composite estimator of monthly unemployment from the Current Population Survey (CPS)). The weights W_i may themselves be complex functions of other weights. (In monthly CPS, for example, the final weights incorporate ratio adjustments for noninterviews and for the first-stage selection of primary sampling units, followed by a two-way raking (iterative proportional adjustment) to national and state population controls.)

The general method of replicate weighting is to assign to each survey case i replicate weights W_{ir}, r=1,...,R. These replicate weights giving rise to replicate estimates, X_r, of the population total

$$X_r = \sum_i W_{ir} X_i \tag{2.2}$$

The variance estimator used with these replicate estimates is typically of the form

$$\mathrm{Var}(X_0) = \sum_r d_r\, (X_r - X_0)^2 \tag{2.3}$$

where d_r, r= 1, ..., R, are independent of the choice of characteristic X but may possibly depend upon the selected sample and upon r. (For many common replication methods, d_r will simply be a constant depending upon the number of replicates R, such as 1/R or (R-1)/R.) Similarly, an estimate of variance for any statistic $S(X_0^{(1)},\ldots,X_0^{(k)})$ that is a function of estimates of population totals, $\hat{X}_0^{(1)},\ldots,\hat{X}_0^{(k)}$, of the form (2.1) and that is continuously differentiable in a neighborhood of the true population totals, is given by

$$\hat{\text{Var}}(S(\hat{X}_0^{(1)},\ldots,\hat{X}_0^{(k)})) = \sum_r d_r \left(S(\hat{X}_r^{(1)},\ldots,\hat{X}_r^{(k)}) - S(\hat{X}_0^{(1)},\ldots,\hat{X}_0^{(k)})\right)^2 \quad (2.4)$$

for the same d_r as (2.3). Substitution of cross-product terms in place of squares in (2.4) estimates the covariance between any two such statistics. In addition, it is a matter of current research to determine situations in which specific replication-based estimates of variance such as (2.4) may give suitable estimates of variance for some specific non-smooth functions S for which linearization fails. (This technical issue is beyond the scope of this paper, however.)

Formulas (2.3) and (2.4) employ variation about the original estimate to measure variance, but alternative variance estimators, such as the complementary half-sample method (e.g. Efron 1982) may also be implemented through replicate weighting schemes. Additionally, other features of replication, such as bias removal of $O(n^{-1})$ for locally quadratic functionals (Fay 1984), may be accomplished in through this device. Replicate weights may be used to implement some variant of the bootstrap replication method (Efron 1982) (although research in this area is still at a developmental level with respect to complex samples.)

The key idea in the replicate weighting representation of replication methods is to associate the replicate weights with the individual survey records, either by placing them in a common computer file or in a file that can readily be matched on a one-to-one basis with the survey characteristics. This approach to the organization of variance computations leads naturally to three distinct phases:

1) Formation of replicate basic weights W_{ir}^* for each survey case based on the survey design and the inverse of probabilities of selection. These weights may be constructed according to a familiar replication method, such as balanced half-sample replication, or according to a more general resampling plan so that their use in (2.2) and (2.3) would lead to estimates of variance (2.4) whose expected value over this randomization of the replication agrees with classical estimators for the variance of the "simple-unbiased" (Horwitz-Thompson) estimator.

2) For each set r of replicate weights W_{ir}^*, recomputation of final weights W_{ir} according to the estimation procedures (ratio estimation, nonresponse adjustments using weighting etc.) applied to the original weights.

3) Use of the replicate weights in (2.2), (2.3), and (2.4) to compute variances for the survey characteristics.

Implementation of replication in this form leads to a highly desirable modularity to the process of estimating variance. The first phase concerns specific aspects of the design that may include stratification, finite population corrections (if needed) etc., while the other two phases need to take no note of these considerations. Similarly, the specific form of the survey estimators to derive the final survey weights need only be reflected in the second phase.

Most importantly, association of replicate weights with the data file permits the computation of design-based estimates of variance for a wide variety of statistics and by potentially a large number of users. The investment in the computation of replicate weights in the first two phases permits economies of scale with respect to computation at the last phase, if a large number of variances are required. Perhaps more important are the economies due to the simplicity in implementing the last phase for simple or highly complex statistics.

The balance of this paper describes a number of applications of this idea and specific features of the computational implementation.

3. THE OCCUPATIONAL CHANGES IN A GENERATION SURVEY; SUBSEQUENT APPLICATIONS AT THE BUREAU OF LABOR STATISTICS

During a seven-year period beginning in the 1950's when CPS had a 230-area design, 40 random replicates were processed through the second stage ratio and composite estimation procedures (U.S. Bureau of the Census 1978). This process took "20 hours of computation each month on a high-speed electronic computer" (U.S. Bureau of the Census 1963) (UNIVAC 1) for a small set of items. To reduce costs, the number of replicates was reduced to 20. With the 1960 redesign and expansion of CPS, the Census Bureau began estimating variances using a procedure proposed by Keyfitz (1957) instead of the replication process.

Once the program for the Keyfitz procedure was developed, the Census Bureau did not compute replication variances until 1973 when the sponsor of a CPS supplement for Occupational Changes in a Generation (OCG) stated it wanted to be able to estimate variances for a variety of analytically complicated statistics that were not specified in advance.

In order to give the sponsor maximum flexibility in variance estimation, the decision was made to use a balanced pseudo-replication procedure. Since the supplement occurred during the phase-in of the 1970 CPS redesign, the development of replicates was complicated by having two different sample designs the 449 PSU design and the 461 PSU design (Dippo 1975). The procedure for constructing the 243 orthogonal one-third pseudo-replicates for this survey is is described by Gurney and Jewett (1975).

Each of the 243 replicates consisting of one-third of the sample was processed through the second stage ratio estimation procedure. Therefore, each record on the sponsor's data file had 243 replicate weights along with a full sample weight. With the OCG survey, an important step was taken with respect to variance estimation: the conscious decision to give the user the flexibility of computing variances as simply as possible. By replicating the weighting process and including the replicate weights on the data file, it is not necessary to specify in advance what item or statistics are to be estimated. This is significant when you consider the fact that variances for only 14 items were computed for CPS during the 1950's. Why? The process of replicating the weighting and estimation was considered as one procedure instead of two.

At about the same time as OCG (1973-75), BLS was revising the Consumer Price Index (CPI) and was laying the ground work necessary to compute variances for the revised index using the balanced half sample replication technique (Weber 1981). One of the major components of the CPI is the set of cost weights or mean expenditures for each market basket-population item (U.S. Department of Labor 1984). The mean expenditures for the index introduced in 1978 were estimated from the 1972-3 Consumer Expenditure Survey. Even though the variances were estimated using random groups and collapsed stratum procedures, it was very time consuming since such a large number of estimates were produced to support both the CPI and the publication of mean expenditures by various demographic, geographic and economic classifications. When BLS began preparing specifications to the Census Bureau for the ongoing CES survey, which began in January 1980, it wanted to insure variances could be produced in a timely manner for a large number of estimates. In addition, the variance estimates should reflect to the extent possible the effects of weighting and sample design. Thus, it was decided to use the 20 half sample replicates set up for the CPI. Although the 20 replicates are an orthogonal set, they do not result in a completely balanced design since there are 58 non-self-representing PSU's in the CPI (Coleman 1974).

In order to reflect the effects of weighting on the variance of the estimates, the entire weighting procedure should be performed independently on each replicate. However, the number of sample units being weighted in a group for CES is only 400 or 200 per half sample. Rather than repeat the collapsing procedure required for the noninterview adjustment independently with each replicate, it was decided to begin the replicate weighting procedure with the second stage ratio estimation. Since the second stage ratio estimation is the only other stage of estimation for CES, basically only the noninterview adjustment procedure is not reflected in the replicate weights. During the second stage ratio estimation procedure, each person 14 years old or older receives 21 different second stage factors, of which 10 are zero. Similarly, each consumer unit receives 21 family weighting factors and 21 final weights. Therefore, each family record produced by the Census Bureau contains as a minimum: one basic weight (the inverse of the probability of selection), one weighting control factor, one noninterview adjustment factor, 21 second stage factors for the principal person, 21 family weighting factors and 21 final weights, for a total of 66 weights or weighting factors. Designing a system to produce 21 times the original number of weights used for estimation does not proportionately inflate the cost: replicate weighting represented only approximately one third of the total development cost for weighting programs and 60-65 percent of the monthly production cost.

BLS loads all of these weights and the data into a RAPID database easily accessible using SAS. Any tabulations or statistics can readily be computed 21 times using each of the 21 final weights. Then it is a simple procedure to compute the average of the sum of the square differences. Even using SAS, which is not the most efficient language with respect to machine time, it only costs about $500 to compute variances for 21 tables which cost $200 to produce without variances. The inclusion of the replicate weights on the data file has allowed BLS to publish estimates of the sampling error for CES simultaneously with the estimates of mean expenditures (U.S. Department of Labor 1983). Not only can the variances of the cost weights for the current CPI revision be easily computed, but for the first time it will be possible to include the variance of the cost weights in the variance of the index. By including the weights as factors for each weighing stage, BLS has been able to look at the effects on the variance of the second stage ratio estimation and family weighting procedures (Dippo 1982). If an analyst at BLS or Census wishes to look at a special tabulation, it is almost as easy to compute the variance as it is to produce the tabulation. The public use tapes issued by BLS for CES contain the final replicate weights and the documentation includes a paragraph on how the user can easily compute variances on any statistic (Jacobs 1982).

4. RECENT APPLICATIONS AT THE CENSUS BUREAU

After the 1973 Occupational Changes in a Generation supplement to CPS, appproximately seven years passed before the next application of replicate weighting at the Census Bureau. This next implementation was to the 1980 Post Enumeration Program (PEP), a sample-based evaluation of coverage in the 1980 census, described in more detail by Cowan and Fay (1984). This survey was based principally on the April and August 1980 samples from the Current Population Survey (CPS) and a sample selected from the 1980 Census (E-sample) in the same counties (primary sampling units (PSU's)) as CPS but selected independently within primary sampling units.

Although use of the CPS design seemingly would have made available existing procedures to estimate sampling variance, the linearization program for CPS variances reflected the design only up to 1975, before extensive changes arising from restratification and supplementation of the sample in many states. A combination of factors - the highly complex design; the objectives of estimating variances for a large number of characteristics at both national, state and specific substate levels; and the complexity of the weighting procedures and of the dual system estimator to estimate the true total population (involving hundreds of cells and covariances between approximately a thousand sample estimates) essentially made linearization infeasible for this purpose. Replication methods implemented through the replicate weighting approach appeared the only reasonable alternative that would both correctly reflect the complexity of design and estimation and also provide estimates of variance for the wide variety of statistics of interest.

One aspect of this application illustrates the flexibility of replication methods to deal with non-standard problems. The CPS design included sets of non-self-representing PSU's drawn in triplets: the first two each as single selections from two strata, and the third by selecting one of the paired strata with equal probability and drawing the third PSU from the selected stratum according to the initial probabilities with replacement. PSU's selected twice were treated as two independent draws in selecting the sample within PSU's, and in estimation. Technical Paper Number 40 (U.S. Bureau of the Census 1978) describes an unbiased estimator of variance for the Horwitz-Thompson estimator under this sampling design: if X_1, X_2, and X_3 represent the totals from the three PSU hits, respectively, where X_1 and X_3 are assumed to come from the same stratum, the estimator is

$$\text{Var} = 1/4((X_1+X_3)/2 - X_2)^2 + 21/16(X_1 - X_3)^2 \quad (4.1)$$

By choosing $d_r = 4/R$ in (2.3), it is possible to generate replicate weights with all W_{ir} positive so that (2.3) has the same expected value (over the randomization used to create the replicate weights) as (4.1). (The usual choice $d_r = 1/R$ could not be used to simulate (4.1) without the potentially undesirable assignment of negative W_{ir}.) The replicate weights mimicked the variation indicated by the first term of the right-hand side of (4.1) for one eighth of the replicates and the second term for seven eighths. More specifically:

	Factors			
X_1	.646 $(=1-2^{1/2}/4)$	1.354 $(=1+2^{1/2}/4)$	1.612 $(=1+(3/8)^{1/2})$	.388
X_2	1.707 $(=1+2^{1/2}/2)$	.293 $(=1-2^{1/2}/2)$	1.000	1.000
X_3	.646	1.354	.388 $(=1-(3/8)^{1/2})$	1.612
Prob	1/16	1/16	7/16	7/16

(The factor (3/8) in the third column is computed as $(21/16)*(8/7)*(1/4)$. These three terms come from (4.1), $(7/8)^{-1}$, and $d_r = 4/R$, respectively.)

The rest of the creation of replicate weights for the design, both for non-self-representing PSU's in restratified states where (4.1) was inappropriate, and for the sample in self-representing PSU's, was similarly performed consistent with $d_r = 4/R$. For example, in instances where a collapsed stratum variance was to be applied to a pair of non-self-representing PSU's, the replicate weights were set up to inflate the weights of one PSU by the factor 1.5 and to multiply the other by .5, randomly from one replicate to another. (Use of 1.5 (=1+1/2) and .5 (=1-1/2) corresponds to $d_r = 4/R$, while factors 2 (=1+1) and 0 (=1-1) are typically used in half-sample replication with $d_r = 1/R$.) Unlike more standard replication methods, each positive initial weight W_i insured positive W_{ir}, which in turn insured that complex functions built from ratios would be defined whenever the function could be computed for the whole sample.

Before its use for the PEP, this approach to variance computation was experimentally applied to the CPS for labor force characteristics in March and April, 1980, using the full CPS design in place at the time. The implementation included all of the complex noninterview and first- and second-stage weighting of CPS, except for composite estimation. The results agreed quite closely with those obtained earlier from the linearized variances for the earlier CPS design, except that the replication-based variance estimates were smaller for some agricultural items, probably reflecting the larger sample size in 1980 in some key agricultural states than for the earlier design upon which the linearization results were based.

The samples for PEP were not independent of each other: a correlation between the E-sample and CPS arose from the use of the same sample of non-self-representing PSU's, and the April and August CPS samples were additionally correlated at the next stage of selection within PSU's. Consequently, an important feature of the application of replicate weights to PEP was to vary the weights in the same manner for dependent components of the sample, thus correctly representing the effect of covariances between the sample designs. For example, in non-self-representing PSU's, the April, August, and E-sample cases each received the same set of multiplying factors in computing replicate basic weights. (This consideration is generally important in the representation of variances for longitudinal surveys with replication.)

The replicate weighting approach enabled the computation of variances for a very large number of individual estimates - literally hundreds of thousands from the PEP. Because of the modularity of the replicate weighting approach, as discussed in the introduction, the validity of the calculations was easily verified.

Replicate weighting was also used to compute variances from the Residential Finance Survey, a follow-on survey sampled from the 1980 census to study characteristics of the financing of housing units. The sample design was similar to other demographic surveys in terms of first and second stages of design, although measures of size for residential properties played a critical role in defining strata and selection rates. Sufficiently large properties were sampled with certainty in self-representing PSU's and therefore did not directly contribute to the design-based variance. (The replicate unbiased weights did not vary for these properties, and the only variation in their final replicate weights came about from the effect of replication of ratio estimation.) Because of the experience gained from the PEP application, implementation of the replicate weighting methodology was easily accomplished.

The English Language Proficiency Study was a relatively small survey, also sampled from the 1980 census. The sample design included the same sets of three PSU's in two strata that appeared in CPS and PEP, but sampling within PSU's and estimation were performed on the basis of unconditional probabilities of selection, rather than the weighting given in CPS and PEP. Consequently, (4.1) was no longer an appropriate variance estimator. Instead, an estimator was derived based upon the Yates-Grundy estimator (Cassel, Sarndal, and Wretman 1977) corrected by an expression that arose from the variation in the number of sampled PSU's in the pairs of strata (most pairs with three but some with two hits). Implementation of such a complex expression would have been difficult to program; instead, replicate weights were constructed consistent with this estimator, thus simplifying the production of variances in this complex situation.

A similar replication approach is currently being implemented for 1984 panels of the Survey of Income and Program Participation (SIPP). Unlike other applications discussed in this section, the second of the three phases described in the introductory section will initially be omitted for simplicity, but, for each case, the ratio of original unbiased weight to the final weight will be applied to all replicate unbiased weights as a rough approximation to the effect of weighting. As time permits, a more complete approach may be taken after the initial publication schedules are met.

The computation for these implementations of replicate weighting at the Census Bureau has been principally in FORTRAN. No completely general programs have been developed, but specific programs, particularly in phase I and to some extent for phase II, have been applied to more than one survey with little or no modification. Thus far, there has been little impetus to produce any sort of general program for phase III, essentially because the actual computation of variance through (2.4) is so straight-forward that there doesn't appear to be any advantage to incorporating this formula in a general-purpose program. In implementing the replicate computations, particularly through phase II, files of weights, weighting keys, and survey identifiers in similar formats have been employed, but phase III has been accomplished by matching a file of replicate final weights to the original data file without changing the original format of the survey data, thus permitting computation of any variance. This design provides the greatest flexibility but has tended to discourage any attempt to produce a general phase III program, since the original data files may substantially differ in format between surveys. More general programs for the first two phases would be desirable, however, and experience with new applications should help to suggest the design of such software.

5. WESTAT'S EXPERIENCE

5.1 Background

Westat's initial efforts to develop a general purpose variance estimation package based upon replication consisted of a program written in FORTRAN which used balanced repeated replication. This program was designed strictly for applications within Westat. It required fixed-format, card-image inputs. It was capable of estimating variances for totals and ratios, and it permitted strata and universes totals to be used for ratio adjustment. Although we found frequent use for the program, it was not designed to be used with more complex statistics, and we thought it would be useful to expand its capabilities.

Under contract to the DOT, we developed an estimation package to permit routine variance estimates for accident characteristics obtained from the National Accident Sampling System (NASS). The NASS, designed in part by Westat, is an anual survey of motor vehicle accidents occurring on our Nation's highways. The design is a nested 75 PSU stratified cluster sample. Interview teams are collecting data throughout the 1984 calendar year in 50 PSU's. At some future point, the remaining 25 PSU's are expected to join the data collection process.

Since the NASS, an on-going survey, is used by many highway safety researchers to examine accident characteristics, it was considered prudent to provide these researchers with a method for gauging the variability of NASS estimates. Westat developed a computer package, referred to as NASSVAR, which is to be embedded in SAS, a commonly used statistical system with the capability of preparing weighted estimates from survey data. The selection of SAS for the programming environment was made by the DOT which uses it frequently for tabulations and analysis of highway safety data.

The use of NASSVAR is accomplished in two steps. First, the replicate definitions are used to compute weights for each case for each replicate sample. These weights are attached to the survey records and are used during the second step, computation of survey estimates, computed statistics and their variance estimates.

5.2 Objectives

One of our primary design objectives was to make a variance estimation program that was easy to use. Several programs have been written to provide variance estimates which have specific card-layout formats or rigid survey data file format requirements. The NASSVAR program is installed in SAS and, as such, is used just as any other procedure in that system. The survey data must be in SAS data file format, a requirement of little consequence to many people using mainframe computers for survey processing.

Another important objective was to allow the user wide latitude in selecting statistics of interest. In this regard, replication methods provide a distinct advantage over the linearization approach. Each new statistic desired must be approximated and new program code written to add the new statistic to a linearization-based approach. The NASSVAR replication approach allows the user to define any SAS provided transformation of weighted survey totals as computed statistics. These statistics are repeatedly computed using the replicate weights after which the variance among the replicated estimates is computed. For example, a log-odds ratio can be estimated and a variance prepared. The user need only specify the four totals involved and a fifth computed statistic defined as the logarithm of the appropriate ratio. Replicated estimates will be computed and a variance among the replicates about the full sample estimate taken.

Another interesting benefit of the NASSVAR program is that it allows the user to define replicates uniquely for each sample design. A user may employ any replication method, for example the Jackknife, Balanced Repeated Replication or any other system for identifying replicates. In fact, the survey researcher interested in comparing various replication methods can prepare several sets of replicate weights and then compute and compare estimates using each. The preparation of replicate weights should include the use of statistical adjustments performed on the full sample if the effect of these adjustments are to be measured by the variance estimation procedure. Adjustments may have been made to account for non-response, undercoverage or other purposes.

The survey totals may have been ratio adjusted or raked to known universe totals. These individual adjustments can be performed easily during the first step of preparation of the replicate weights. During the development of NASSVAR we programmed a simple set of matrix operations which can be performed by the PROC MATRIX procedure of SAS. These operations will compute the necessary ratio adjusted replicate weights while allowing for a series of ratio-adjustment steps.

5.3 Cost

Our experience using NASSVAR with a number of studies indicates that the computer resources required can be predicted very accurately knowing the number of survey records, number of replicates and number of survey variables involved in the computation. The costs are linearly related to the product of these three numbers. A regression on a number of computer runs suggests the following predictor of CPU time on an IBM 370

CPU seconds $= 4.4 + 2.1*10^{-5}$ * replicates * records * estimates

5.4 Sample Design Restrictions

The basic requirement is for a stratified sample of PSU's at the first stage. To date, only balanced repeated replicates have been used requiring two selections per strata. In this case if only one selection is taken to maximize the effect of stratification, the strata are paired into pseudo-strata so that half sample may be taken. To prepare estimates for designs with self-representing PSU's, the within strata design must be taken into account to create two pseudo-strata which describe the sampling error from that step of the design. For the NASS design we used alternative sampling weeks to define two paired strata. In other designs, when the sampled cases were selected in a systematic manner, odd vs. even case number was used to define the two pseudo-PSU's. Other replication methods, such as the Jackknife, would not require two selections per first stage stratum.

5.5 Current Activities

Westat is currently completing an interface program for an analyst wishing to perform multi-way contingency table analysis on survey data. The program uses CPLX written by Robert Fay, one of the authors. CPLX prepares a contingency table analysis of survey data using the method of maximum likelihood. The purpose of the interface, written in the WYLBUR command language, is to interrogate the user about the table model and hypotheses to be tested. The program then generates an input file required by CPLX and reads the survey data file which was previously written in SAS format.

In addition, Westat has begun the extension of the NASSVAR program to include a multiple regression model. As with the contingency table approach, replicated analyses will be performed and regression coefficients estimated for each replicate. The variance of the coefficients will be estimated by taking the sample variance of the replicates about the full sample estimate. This enhancement should prove valuable to the survey researcher.

REFERENCES

Cassel, C.-M., Sarndal, C.-E., and Wretman, J.H. (1977), Foundations of Inference in Survey Sampling, New York: John Wiley.

Coleman, J. (1974), "The Continuing Consumer Expenditure Survey," unpublished memorandum, U.S. Bureau of Labor Statistics, U.S. Department of Labor.

Diaconis, P. and Efron, B. (1983), "Computer-Intensive Methods in Statistics," Scientific American, 248, 116-130.

Dippo, C.S. (1975), "Variance Estimation for Occupational Changes in a Generation (OCG) Survey," unpublished memorandum, U.S. Bureau of the Census.

______ (1982), "Preliminary Report on 1980 Diary Weighting," unpublished memorandum, U.S. Department of Labor, U.S. Bureau of Labor Statistics, March 1, 1982.

Efron, B. (1982), The Jackknife, the Bootstrap and Other Resampling Plans, Philadelphia, PA: Society for Industrial and Applied Mathematics.

Fay, R.E. (1984), "Some Properties of Estimates of Variance Based on Replication Methods," for presentation to the Annual Meetings of the American Statistical Association, Section on Survey Research Methods.

Gurney, M. and Jewitt, R. (1975), "Constructing Orthogonal Replication for Variance Estimation," Journal of the American Statistical Association, 70, 819-821.

Jacobs, C. (1982), "Estimation Procedures for the CES Dairy Public Use Tape," memorandum dated July 16, 1982, U.S. Department of Labor, U.S. Bureau of Labor Statistics.

Keyfitz, N. (1957), "Estimates of Sampling Variance Where Two Units Are Selected from Each Stratum," Journal of the American Statistical Association, 52, 503-510.

U.S. Bureau of the Census (1963), The Current Population Survey: A Report on Methodology, Technical Paper No. 7, by Joseph Steinberg, Thomas B. Jabine, and Leon Pritzker, Washington, D.C.: U.S. Government Printing Office.

______ (1978), The Current Population Survey: Design and Methodology, Technical Paper No. 40, Washington, D.C.: U.S. Bureau of the Census.

U.S. Department of Labor, Bureau of Labor Statistics (1983), Consumer Expenditure Survey: Diary Survey, 1980-1981, Bulletin 2173.

______ (1984), 1984 BLS Handbook of Methods, Vol II, The Consumer Price Index, Bulletin 2134-2.

Weber, W. (1981), "A System of Variance Estimation for the U.S. Consumer Price Index," Proceedings of the Section on Survey Research Methods, American Statistical Association, Washington, D.C.