Journal of Official Statistics
Vol. 9. No. 4, 1993, pp. 795–814

An Evaluation of the Use of Personal Computers for Variance Estimation with Complex Survey Data

Barbara Lepidus Carlson, Ayah E. Johnson, and Steven B. Cohen[1]

Abstract: Over the past decade, software with the capability of handling data resulting from surveys with complex designs has become more readily available, primarily for use on mainframe computers. More recently, a few of these have been released in versions designed for the personal computer (PC). As the speed and capacity of PCs increase, there is the potential for more widespread use of this specialized software.

In this paper, the authors present the results of comparisons between two of these PC software packages, SUDAAN and PC CARP, and their mainframe counterparts with respect to computer run-time, programming difficulty, special features and flaws, and data preparation issues, as well as space and memory requirements for the PCs and cost considerations for the mainframe. Analyses are limited to computations of variance estimates, using Taylor linearization, for weighted proportions and means, and are carried out on files from the 1987 National Medical Expenditure Survey, which has a stratified multi-stage area probability design.

Key words: Statistical software; SUDAAN; PC CARP.

1. Introduction

Many national surveys have sample designs that deviate from simple random sampling. Stratification is often considered to increase the precision of survey estimates. Clustering is frequently used to make the field work of the survey more efficient. In addition, when greater representation of certain policy–relevant subgroups is necessary, disproportionate sampling is often used. Sampling weights are calculated to reflect the unequal probabilities of selection.

Mainframe computers have been the primary resource used to support federal research and analysis. Since standard statistical computing packages, such as SAS and SPSS, assume simple random sampling, any variance estimates arising from them may not reflect the actual variance achieved by adoption of a more complex design. Specialized software which accounts for complex survey designs when estimating variances has existed for about a decade, but primarily for use on mainframe computers.

[1] Center for General Health Services Intramural Research, Agency for Health Care Policy and Research, Executive Office Center, Suite 500, 2101 E. Jefferson Street, Rockville, MD 20852, U.S.A.
Acknowledgement: The views expressed in this paper are those of the authors and no official endorsement by the Department of Health and Human Services or the Agency for Health Care Policy and Research is intended or should be inferred. The authors would like to thank the reviewers for their helpful comments.

With the increased prevalence of personal computers (PCs) and their increasing capacity and speed, the idea of analyzing survey data files on PCs becomes more plausible. Using PCs rather than mainframe computers has many potential benefits as well as costs. The benefits generally relate to cost savings, but freedom from some aspects of mainframe computing also makes PCs more attractive. While actual execution time on the mainframe may be substantially less than on a PC, the time from submission of a batch job to the receipt of the printout is generally much longer on a mainframe if the mainframe is off-site. Mainframes often operate on a time-sharing basis, which may mean waiting in an execution or print queue, or both, during a busy period. All mainframe computers have down-times, some more often than others. Due to prohibitive expense, large jobs are often submitted for execution during a discount time, usually overnight or over the weekend, which substantially slows down the entire process. Accessible storage of software as well as data on a mainframe is a daily expense.

One generally has immediate access to printouts when using a PC. PC packages often have an interactive or menu format, rather than a batch format, which can make a package easier to learn and use. Furthermore, the freedom from Job Control Language, used to inform the mainframe operating system how to process the job, is an attraction of the PC. There is usually an option with PC packages to output tables and other results from an analysis into a separate text file, which can be quite helpful later for creating tables without retyping the numbers. Finally, distribution and installation of software is significantly easier on a PC than on a mainframe. Installing software via diskette can be as easy as a copy command, whereas mainframe tape installation requires expert assistance.

Costs associated with using PCs relate primarily to the run-time issue, as well as space and memory constraints. Obtaining computing equipment well-suited for statistical analysis can become quite expensive, since it must have sufficient memory, disk space, speed, and often a mathematical coprocessor. Unless one has a memory manager which allows for several tasks to be performed simultaneously, a PC can be "tied up" while running a lengthy analysis or downloading a file. The need to download data files from the mainframe, along with its potential for introducing transmission errors into the data base, is also a consideration. The downloading process may need to be repeated if the files or analysis plans change.

Since large data bases require significant resources in order to be analyzed in a timely and efficient manner, under what circumstances will the PC versions of variance estimation programs be useful in keeping computing costs down without sacrificing the efficiency normally associated with the power of a mainframe computer? Two of the more frequently-used mainframe packages for the analysis of complex survey data now have PC counterparts: SUDAAN and PC CARP. In this paper, each of these PC programs is evaluated relative to its mainframe version, and the two PC programs are compared to each other. Features available in these packages as well as issues related to the actual implementation of the programs, including data preparation steps, number of programming statements, and time and cost issues, are examined using two data sets from the 1987 National Medical Expenditure Survey (Edwards and Berlin 1989; Cohen, DiGaetano, and Waksberg 1991), which has a complex sample design. Programming statements and samples from resulting output are available from the authors.

2. Background

Most of the commonly-used statistical computing packages (SAS, SPSS, BMDP) assume data were obtained from a simple random sample; that is, the observations are independent and identically distributed. When data have been collected from a survey which has a complex sample design, the simple random sample assumption can often lead to an underestimate of the variance, which can therefore lead to artificially small confidence intervals and anticonservative hypothesis testing; i.e., rejecting the null hypothesis when it is in fact true, more often than indicated by the nominal Type I error level.

A few different statistical strategies have been developed to address this issue. Among them are: a first-order Taylor series expansion of the variance equation; a balanced-repeated replication method (BRR); and the Jackknife approach (Wolter 1985). Several software packages have been developed which incorporate one or more of these strategies into their variance calculations.

The current evaluation focuses only on software packages which currently have PC counterparts: SESUDAAN, SUDAAN, and SUPER CARP on the mainframe, SUDAAN and PC CARP on the PC. All of these packages use a Taylor series approximation to compute variance estimates. Other programs which are designed to analyze data from complex surveys exist (OSIRIS PSALMS and OSIRIS REPERR from University of Michigan, WESVAR, WESREG, and WESLOG from Westat, Inc., and HESBRR from NCHS), but have no PC counterparts to date, and are therefore not pertinent to the subject matter of this paper. These mainframe packages have been evaluated elsewhere (Cohen, Burt, and Jones 1986; Cohen, Xanthopoulos, and Jones 1988).

SUDAAN (Shah, La Vange, Barnwell, Killinger, and Wheeless 1989) and SESUDAAN (Shah 1981) are two programs designed to estimate variances from complex survey designs and have been developed by the Research Triangle Institute (RTI). Among many other capabilities, these programs can compute weighted means and proportions and their corresponding standard errors (Table 1). SESUDAAN is a mainframe package only, while SUDAAN has PC, mainframe and VAX/VMS versions. It should be noted that SUDAAN on the mainframe was still in test mode and not yet available to the public as of the time of this evaluation. Any comments regarding the mainframe SUDAAN should be regarded in this context. SUDAAN will accept both SAS and text data files. SESUDAAN will accept only the SAS data format.

RTI has produced a family of programs to compute variance estimates, including regression and logistic regression packages. While most of these programs were written in SAS language (SAS Institute, Inc. 1985), the newest of them, SUDAAN, is written in "C." SUDAAN incorporates the features of SESUDAAN, SURREGR (regression), RTILOGIT (logistic regression), RATIOEST (ratio estimation package), and RTIFREQS (frequencies), and has many improvements over these older versions. Although RTI refers to both the mainframe and PC versions as "SUDAAN," for purposes of clarity, the PC version will henceforth be referred to as "PC SUDAAN" in this paper.

It should be noted that the earlier complex survey design software developed by RTI, which includes SESUDAAN and SURREGR, was tailored to applications

Table 1. Comparison of software capabilities

	SESUDAAN	SUDAAN	SUPER CARP	PC SUDAAN	PC CARP
Host system	SAS	C	FORTRAN G	C	FORTRAN G
Accepted data file format	SAS	SAS or text file	text file	SAS (PC) or text file	text file
Mode of execution	Batch	Batch	Batch	Batch	Interactive (menu-driven)
Automatic collapsing of strata[1]	No	No	Yes	No	Yes
Labeling of categorical variable values	Yes	Yes	No	Yes	No
Missing data	· or blank or LE 0 (categ.)	· or blank or LE 0 (categ.)	Won't accept (will do "screening")	· or blank or LE 0 (categ.)	Won't accept (requires imputation)
Max. number of variables	(none)	(none)	50 per analysis (75 variable max.)	(none)	50
Chi-square test of independence	No	Yes	Yes (but no *p*-value)	Yes	Yes (but no *p*-value)
Types of estimates (and standard errors) available:					
Totals	Yes	Yes	Yes	Yes	Yes
Means	Yes	Yes	Yes	Yes	Yes
Proportions	Yes	Yes	Yes	Yes	Yes
Quantiles	No	Yes	No	Yes	Yes
Ratios	Yes	Yes	Yes	Yes	Yes
Differences	Yes	Yes	Yes	Yes	Yes
Cross-tabs	Yes	Yes	Yes	Yes	Yes
Design effects	Yes	Yes	No	Yes	Yes
Regressions	(SURREGR)	Yes	Yes	Yes	Yes

[1]When less than two PSUs exist for a stratum.

that coincide with the analytical requirements of the National Medical Expenditure Survey (see Section 3.3). The new analytical capabilities that are available in SUDAAN and PC SUDAAN have been added as a consequence of funding provided by a Public Health Service Task Force. More specifically, funding support was provided for new software development to meet analytical requirements for the analysis of complex survey data, primarily obtained from health care surveys sponsored by the Department of Health and Human Services.

SUPER CARP (Hidiroglou, Fuller, and Hickman 1980) and PC CARP (Fuller, Kennedy, Schnell, Sullivan, and Park 1989) are products of the Statistical Laboratory at Iowa State University. SUPER CARP is a mainframe package, the latest version of which is approximately ten years old. PC CARP, its PC counterpart, is relatively recent, and has many improvements over its mainframe parent. These programs are written in FORTRAN G. Two supplemental programs, logistic regression and post-stratification, are also available. SUPER CARP and PC CARP will accept only text data files. Like the RTI programs, these two packages have many statistical capabilities other than those being evaluated here (Plewis 1989).

SUPER CARP and PC CARP development was partly supported by the U.S. Department of Agriculture (Soil Conservation Service), Bureau of the Census (International Statistical Programs Center and Methodology Division), and by Doane Agricultural Service (St. Louis).

3. Methods

The computer packages were evaluated with respect to efficiency, accuracy, and ease of use on both a mainframe and a personal computer, since each package has a version for both environments. For each package, weighted means and weighted proportions were estimated for the same set of variables on each of two similar data sets. Along with these estimates, standard errors and design effects[2] were also computed. The evaluation of these software packages is done by examining the following features:

- Comparison of features and flaws
  - SESUDAAN vs. PC SUDAAN
  - SUPER CARP vs. PC CARP
  - PC SUDAAN vs. PC CARP
  - Handling of missing data
- Programming effort
- Execution time
- Computing costs
- Computational accuracy
- Tests of association
- Quality of documentation.

[2] A design effect is the ratio of the variance of the statistic under the actual design divided by the variance that would have been obtained from a simple random sample (Cox and Cohen 1985).

3.1. *Procedures*

All of the programs being evaluated require that the data be sorted by the nesting variables. SAS data sets sorted by stratum and PSU were created on the mainframe, keeping only the variables and observations relevant to the evaluation, and converting missing values (originally coded as negative numbers) to dots, the SAS convention for missing numeric data. (See Figure 1.) These SAS data files were used for SESUDAAN and SUDAAN runs. Text files were then created from these SAS files using SAS and then WYLBUR, the mainframe's online interactive text editor, to change dots to blanks for missing data (the FORTRAN convention for missing data), and were used for SUDAAN and SUPER CARP runs. SUDAAN will accept both text and SAS data files. The text files were downloaded to the PC using MS-DOS Kermit software. These text files were then used on the PC for PC SUDAAN and PC CARP runs. The files were input into the PC version of SAS (SAS Institute, Inc. 1988) and then used for PC SUDAAN runs once again. When downloading a data file from the mainframe for use in PC SUDAAN, the data were first converted to text format and then read into the PC version of SAS after the download was complete. (A direct SAS download via its Micro-to-Host Link is possible, but complex.) One must keep in mind that SAS on the PC requires a substantial amount of hard disk space, and it may not be feasible for many to have this capability on their PCs.

In addition to the data file and the program code, several other files were necessary in order to run SUDAAN in both mainframe and PC environments. For

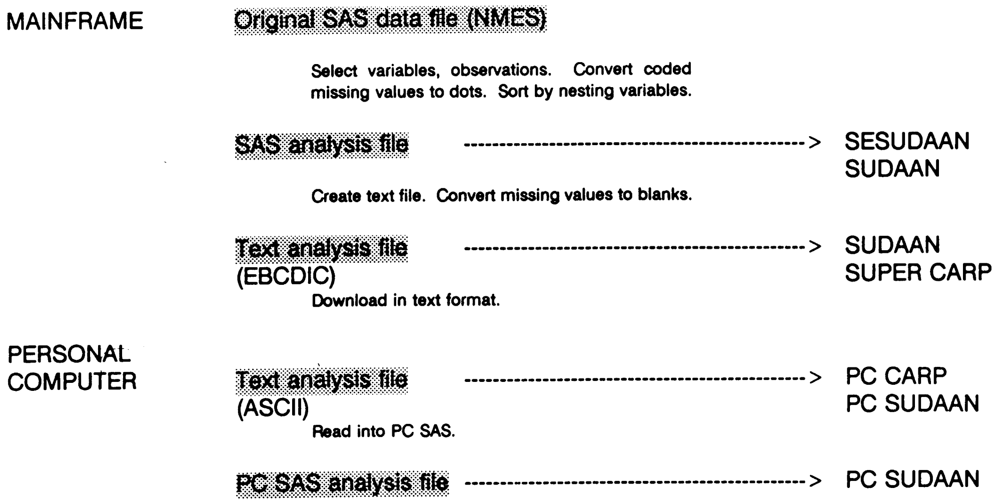

Fig. 1. Steps in processing of data files

SUDAAN, when using a SAS data file, one need only create in addition a "LEVEL" file, which contains labels for different values of any or all categorical variables, as would appear in a SAS PROC FORMAT value statement. This file is optional. When using a text data file in SUDAAN, this "LEVEL" file is also optional, but an "LBL" file and a "LABELS" file are always required, regardless of whether there are categorical variables, listing the variables in the order in which they appear on the data set, and variable descriptors (length, type, file name), respectively. These auxiliary files were downloaded to the PC as well, using Kermit. No auxiliary files were necessary to run SUPER CARP or PC CARP. For all packages used, two subsampling levels were specified. The CARP packages allow for only two levels of subsampling, while the RTI packages allow for any number of nesting levels. The authors did not use a finite population correction factor (fpc), since there generally is a negligible gain from this specification with the National Medical Expenditure Survey data. For SESUDAAN and SUDAAN, the fpc is specified through counts, whereas for SUPER CARP and PC CARP, it is specified as rates.

Programming effort was measured by the number of statements required to run the program. For the SESUDAAN programs, a statement was defined as ending with a semicolon. For SUDAAN and PC SUDAAN programs, the number of statements was defined as the number of program statements ending with a semicolon (excluding Title statements) plus the number of lines required in all of the necessary auxiliary text files. For SUPER CARP, the number of statements was equal to the number of lines in the program, excluding the extra lines needed to enumerate the variable names. Job Control Language, SAS options statements, and Title statements were excluded from all measurements on the mainframe. Since PC CARP was not run from a batch program, the number of responses to PC CARP prompts was recorded instead.

When writing the programs to execute the packages being evaluated, an attempt was made to minimize the number of steps needed to execute the program and to make the runs on the various software packages as similar as possible. Although SUPER CARP and PC CARP do not allow for formats (i.e., labeling of variable response categories), use was made of this capability for the SUDAAN and SESUDAAN runs. In SUPER CARP and PC CARP, extra programming steps became necessary due to the way in which the program handles missing data (described below in more detail).

Execution times and computing costs, two of the outcomes of interest, were automatically computed and recorded on the printed output from the mainframe runs. A precise execution time for PC CARP is difficult to obtain, as well as inappropriate, due to the different execution modes: the interactive nature of PC CARP (and its dependence on user response and key-entry speed) versus the batch nature of PC SUDAAN. Therefore, for all of the PC runs, approximate run-times from the first keystroke to the final result were recorded. When several passes were required to complete the analysis (due to memory constraints or variable limitations), the run-times reflect the sum of the discrete times of each of the runs. Computational accuracy was evaluated by examining the output from the various programs and determining at which decimal point discrepancies between estimates began to occur.

SESUDAAN required two procedure executions (PROC SESUDAAN), one for computing means and one for the proportions. Similarly, SUDAAN required a PROC DESCRIPT procedure for means and a PROC CROSSTAB procedure for proportions. For SESUDAAN, a dummy variable uniformly equal to 1 had to be created which was used as a "report" (subgroup) variable when computing overall means. PC SUDAAN ran out of memory when computing the proportions; therefore, the means of the 11 continuous variables were calculated in one run, the proportions of the first 12 categorical variables were the second run, and the proportions of the last 12 categorical (dichotomous) variables were the third run. Note that, for the dichotomous variables, mean estimates could have been made, rather than proportional estimates. Variables are described further in Section 3.3.2.

SUPER CARP and PC CARP each required two runs due to the limitation of a maximum of 50 variables that can be read into the program. One run estimated the means, and another the proportions. In order to estimate means of variables with no missing values, the packages require that a ratio with an all-ones variable in the denominator be estimated. This is done to sum the weights in the denominator of the estimate, the numerator being the weighted sum of the variable of interest. (In PC CARP, a mean can also be computed using the "Subpopulation Means" option with the all-ones variable as the classification variable.) For proportions of variables with no missing values, "Subpopulation Proportions" were estimated, with the all-ones variable being the "dependent variable" and the variable of interest being the "classification variable."

SUPER CARP and PC CARP automatically treated missing data, represented by blanks, as zeros. They will not accept missing data represented by dots. In order to get around this problem, missing value indicator variables had to be created for each variable containing missing data, more than doubling the size of the text files. In both SUPER CARP and PC CARP, these indicator variables were treated as subgroup variables. Weighted means were then computed using the "Subpopulation Means" option with the missing value indicator variable being the classification variable, with one of the two "subgroups" yielding a mean of zero.

It was more difficult to obtain variance estimates for weighted proportions with missing values using the CARP programs. On the mainframe, the "Subpopulation Proportions" option was used, with the "dependent variable" being specified by the missing value indicator variable. In these instances, the "maximum number of categories" for the analysis was specified as twice the number of categories in the variable of interest. On the PC, this situation was more directly dealt with using the newly-available "Two-way Table" option, using the indicator variable and the variable of interest in the cross-tabulation. One must then count the missing value as a category when specifying the number of categories. It is important to note that we ordinarily correct for missing data prior to making estimates. An alternative in the CARP environment would be the elimination of observations with any missing values if an exploratory analysis wished to determine the mean estimate and associated variance for respondents.

3.2. *Computing environment*

The mainframe computer used is an IBM 3090 Model 300J located at the National Institutes of Health in Bethesda, Maryland. It runs under the OS/MVS/ESA operating system. The personal computer used is an AGI 3000D (a 386 IBM clone) with an 80386 Processor, 80387 Numeric Co-Processor, 40 mb hard disk, and 3 mb extended RAM. It runs under the MS-DOS (version 3.30) operating system.

3.3. *Data files*

3.3.1. The survey data

The 1987 National Medical Expenditure Survey (NMES)[3] included two distinct household surveys. The first is a national probability sample of the civilian, noninstitutionalized U.S. population. The second is a survey of the American Indian and Alaskan Native population living on or near reservations and eligible for services provided or supported by the Indian Health Service (IHS). Both household survey components were designed to provide statistically unbiased national estimates of health care utilization, expenditures, and access to care, and health insurance coverage for their respective target populations for calendar year 1987. To provide focused estimates of subpopulations of particular policy concern, the Household Survey oversampled the elderly, those with difficulties in performing activities of daily living, poor and low-income families, and the black and Hispanic minorities.

The Household Survey (HHS) and Survey of American Indians and Alaska Natives (SAIAN) sample designs can be characterized as stratified multi-stage area probability designs with three stages of sample selection:

(1) selection of PSUs (counties or groups of contiguous counties)

(2) selection of area segments within PSUs;

(3) selection and screening of dwelling units within segments. The total round one HHS sample comprised 36,400 individuals in roughly 15,000 households. The round one SAIAN sample comprised roughly 6,500 people in approximately 2,000 households.

3.3.2. Variables

Each of the two data sets (HHS and SAIAN) consisted of stratum and primary sampling unit (PSU) indicators, a sampling weight, and 35 variables for which estimates of weighted means or proportions, as appropriate, were computed. There are 11 continuous variables and 24 categorical variables

[3]Sponsored by the Agency for Health Care Policy and Research (AHCPR), formerly the National Center for Health Services Research and Health Care Technology Assessment (NCHSR).

with up to seven categories each, half of which are dichotomous. Missing data existed in the file for most of the variables, due to nonresponse (at most 10%) as well as intentional skip patterns in the questionnaire. The missing value indicator variables added an additional 22 variables to the HHS file and 21 variables to the SAIAN file. The inclusion of missing data enabled evaluation of how each software package handled these cases, although in reality missing values are generally imputed or adjusted for in making national estimates.

The variables chosen for this analysis pertained to sociodemographic data, health status, health maintenance, a description of the medical and dental care environment, and access to health care. These variables were chosen such that roughly one-third of them were dichotomous, one-third categorical (more than two categories), and one-third continuous.

The NMES Household Survey was larger than the SAIAN survey because of the targeted population being sampled and the stricter precision requirements associated with the HHS. Moreover, the number of strata and PSUs was larger for the NMES HHS than for the NMES SAIAN. In order to test the effect on computer time of this dimension, and to reduce the number of observations to a suitable size, a subset of the HHS population was selected (only non-whites were included). Furthermore, the SAIAN and HHS files were limited to round-one respondents. For the software evaluations, responses to the "Access to Care" questionnaire supplement, the "Self-Administered[4] Health Questionnaire," and the main questionnaire for demographic characteristics were used. Those who did not respond to any of these questionnaires were excluded from consideration. Complete documentation on questionnaires and data collection methods is presented in Edwards and Berlin (1989).

[4]Administered by trained interviewers for the SAIAN.

3.3.3. File size

The SAIAN data set had 5,584 observations, and the HHS data subsample had 8,310 observations. For SAIAN, there were 11 strata, with two PSUs per stratum. Each PSU had between 53 and 788 observations. For HHS, there were 71 strata, with two PSUs per stratum, each PSU having between two and 298 observations. Three strata had to be collapsed with adjoining strata due to empty PSUs for the HHS subsample chosen.

The SAIAN text file had a logical record length of 77, and took up about 450 kb. The SAS version of the SAIAN data file took about 950 kb of space on the mainframe, 1.7 mb of space on the PC. The HHS text file had a logical record length of 71, and took up approximately 610 kb of space. The SAS version of the HHS data file took up 1.4 mb on the mainframe, and 2.6 mb on the PC.

4. Results

4.1. Comparison of features and flaws

4.1.1. SESUDAAN vs. PC SUDAAN

Since the qualities of SUDAAN and PC SUDAAN (other than efficiency) are virtually identical, PC SUDAAN is compared here to its mainframe counterpart SESUDAAN only. PC SUDAAN has many advantages over its predecessor SESUDAAN, although some flexibility was lost in the transition. With the change from SAS to C, some options available under SESUDAAN disappeared, making it more difficult to describe the data as well as to modify variables and re-run the program. SESUDAAN had the benefit of being a procedure within a SAS program. Data could be easily manipulated within the same program using SAS DATA steps, whereas the data set running under PC SUDAAN has to be a finalized data set. Categorical variable value labels could previously be specified in the same program.

SESUDAAN was more limited, however, in many other respects. When proportions were being estimated, the standard output showed the unweighted denominator but not the numerator. The weighted numerator was available, but had to be calculated from the given information. In addition, one must divide variables into "report" versus "analysis" variables for each run; it is not necessary to make this distinction with PC SUDAAN. PC SUDAAN's new features include Chi-square testing and weighted quantiles. Row, column, and overall percentages are also available. It allows for post-stratified estimates and has much more flexibility with respect to the specification of the sample design. In addition to the "with-replacement" design, the first stage can be specified as "without-replacement" for either a simple random sample or unequal probability of selection design. Later stages can be specified as "with-" or "without-replacement." PC SUDAAN also produces marginal totals automatically, something which SESUDAAN had to be "tricked" into doing, by creating an all-ones report variable.

4.1.2. SUPER CARP vs. PC CARP

PC CARP has several improvements over its predecessor, SUPER CARP. In SUPER CARP, programming instructions must be specified in particular columns, in a somewhat scattered set pattern. The order of the rows of code is not at all flexible. In PC CARP, that is not an issue, due to the menu-driven mode of specification. In addition, PC CARP computes design effects and quantiles, and automatically prints relative standard errors ("C.V.s"), none of which SUPER CARP produced. Both PC CARP and SUPER CARP offer tests of independence for categorical variables. While not the case with SUPER CARP, PC CARP's categorical variable responses are now displayed in numeric order (for one-way tabulations only). SUPER CARP's system requires an explicit description of the input format (in FORTRAN notation), while PC CARP also has a list format option.

4.1.3. PC SUDAAN vs. PC CARP

PC CARP and PC SUDAAN are comparable in their space and memory requirements. Version 5.02 of the PC SUDAAN software takes up roughly 790 kb of disk space, and requires 640 kb RAM. It can run on any IBM-compatible PC. The PC CARP software takes up roughly 530 kb, including PRE CARP, and requires 450 kb RAM. It can run on an IBM-compatible machine but must have a mathematical coprocessor.

PC CARP software costs 300 USD, with supplements costing 50 USD each, and the version of PC SUDAAN used cost 250 USD. Current versions of PC SUDAAN cost about 1,000 USD. The up-front cost of hardware accommodating such software packages can be over 5,000 USD, assuming no hardware already existed. These costs are per person and, for data handled by many analysts, could imply large investments in PC hardware and software.

An advantage of the PC CARP software over the PC SUDAAN software is the ability to automatically collapse strata when they contain less than two PSUs. In addition, it has an on-line help facility. Disadvantages include the limitation of only being able to read in a text file, which limits making changes in the data, requiring importation back into SAS or a data base manager for modifications, then exportation back to a text file, as well as having to eliminate missing value dots which are output to the text file from a SAS file. In addition, there are limitations to the number of variables that can be read into the program.

The PC CARP software does not accommodate categorical variable value labeling, and for straightforward mean computation, either a ratio with the denominator being the sum of the sampling weights, or a subpopulation mean with an all-ones variable for the subpopulation must be computed. Outputting usable data files (not just text files) comprising the computed estimates is not an option with the PC CARP software. Although significance tests are often available, the test statistic is printed without a *p*-value. Numeric results are expressed in scientific (exponential) notation. Most statements in PC SUDAAN can be in virtually any order (aside from printing specifications), whereas PC CARP is menu-driven.

Other than these observations, the main difference between the two PC packages is in their mode of execution. Having to specify multiple analyses via the interactive mode of PC CARP became tedious, since the 35 analyses run during our evaluation had to be specified individually with separate runs beginning with the Analysis Specification screen. Each estimate, its two component variables (including either the all-ones variable or a missing value indicator variable), and their number of categories had to be specified separately and set to run separately.

One potential improvement to PC CARP would be a batch mode option, which would be especially helpful for large numbers of estimates, when one tends to write out in advance the types of details which correspond to PC CARP's prompts. Currently, parts of the required entry can be stored in a separate file and read in: variable names, variable input formats, and stratum sampling rates. Two problems with PC CARP's interactive mode are that keyboard entry errors are not always easily changeable (although one can sometimes go back to a previous screen using a go-back option), and it is also difficult to determine when PC CARP is actually executing, rather than waiting for your response to the previous prompt, since the screen does not change.

4.1.4. Handling of missing data

How these packages deal with missing data turned out to be one of the most difficult issues. SESUDAAN and SUDAAN treat blanks and dots as missing values for continuous variables in much the same way that SAS does; i.e., any variable being used for a particular computation that has missing data will be excluded for that computation, unless one specifies an option otherwise. In a two-way table where there are missing values for one or both of the classification variables, the observations with any missing value for either variable are excluded from all marginal totals and proportions. By the same token, in regression analysis, any observation with a missing value for any of the variables in the model is deleted from the analysis. Values of categorical variables outside the specified range were also treated as missing in both of these packages, including values less than or equal to zero. For data with missing values coded as negative numbers, as was the case with NMES, one must be sure to convert categorical variables to positive integers, if one needs to obtain an estimate of the representation of item nonresponse, and these negative values clearly need to be converted to dots for continuous variables.

As explained in the Methods Section,

SUPER CARP and PC CARP are quite limited in how they deal with missing data. Basically, they do not accept missing data, and when encountered, a missing value is treated as a zero (unbeknownst to the user). In SUPER CARP, one can specify screening operations, which will screen out observations based on equality or inequality statements. This would have worked fine had the negatively-coded missing values not been previously converted to dots (to accommodate SAS) and then to blanks (to accommodate FORTRAN). But SUPER CARP does not allow the specification of deleting values equal to "blank." Furthermore, the major problem with the screening option in SUPER CARP is that whole records are deleted, rather than excluding the one value needed for the estimate; i.e., it is a list-wise deletion. When multiple variables and multiple estimates are being considered, one does not always want a whole record deleted from the computation of all estimates simply because one of its variables has a missing value.

PC CARP does not have a screening procedure, but comes with a hot-deck imputation program, called PRE CARP, which can be run prior to the PC CARP run. Unfortunately, a hot-deck imputation is not always desired, especially in this evaluative situation, where comparability with other programs was important. If one chooses not to use the PRE CARP, PC CARP will treat missing values as zeros.

4.2. *Programming effort*

As described in the Methods Section, programming effort was measured in terms of number of programming statements required. Disregarding Job Control Language for mainframe packages, and statements used to label categorical variable responses, SUDAAN and PC SUDAAN using a SAS data file and SESUDAAN required the fewest programming statements (see Table 2). Using a text data file in SUDAAN requires several extra files to help describe the data, which increases the programming effort at least six-fold, both on the mainframe and the PC. SUPER CARP requires a large number of state-

Table 2. Programming statements[1] required for 35 estimates[2]

	Number of statements (excluding label-related statements)	Number of statements related to categorical variable value-labeling
SESUDAAN	15	16
SUDAAN and PC SUDAAN (using SAS data file)	11	89
SUDAAN and PC SUDAAN (using text data file)	87	89
SUPER CARP	82	N/A
PC CARP	353[3]	N/A

[1]Excluding JCL for mainframe runs.
[2]"Estimate" here refers to one mean (for a continuous variable) or a set of proportions (for all values of a categorical variable).
[3]Number of Keyboard responses to prompts.

N.B. The SAIAN and HHS data sets required the same number of programming statements.

ments since every estimate requires an extra set of specification lines. Since PC CARP does not have a batch mode, the number of necessary responses to prompts for the estimates run were recorded instead. This resulted in a significantly more extensive "programming effort" for PC CARP, with 353 "statements" versus the SUDAAN program (SAS data) with 11 programming statements.

Labeling values of categorical variables becomes a bit more cumbersome when one leaves the SAS environment of SESUDAAN, and its PROC FORMAT, for the SUDAAN environment, which requires an extra text file with one line per value. The increase is exaggerated in the Table, since multiple lines can be counted as one "statement" in a SAS program, but each line counted as a statement in the text file. Also, the dichotomous variable codes did not actually need to be labeled. Neither SUPER CARP nor PC CARP have the ability to label values of categorical variables.

4.3. Execution time

CPU time on a mainframe is system-specific, as is run-time on a PC. Even for a given PC, hard disk specifications and software caching can alter times. Although the times reported here can vary from system to system, they are presented in order to give a sense of the magnitude of difference between the packages. The IBM mainframe at NIH is considered comparable to other IBM mainframe systems with respect to CPU time.

On the mainframe, CPU time varied widely among the packages. As one can see from Table 3, SESUDAAN required the least time (7 seconds for the SAIAN data, 11 seconds for the HHS data), and SUDAAN (using a text data file) took the

Table 3. CPU time (in seconds) for mainframe packages

	SESUDAAN	SUDAAN (using SAS data file)	SUDAAN (using text data file)	SUPER CARP
CPU time				
SAIAN data set	7.22	52.04	71.21	50.71
HHS data set	10.58	85.07	114.81	81.80
Comparison of CPU time to SESUDAAN (ratio)				
SAIAN data set	(1.0)	7.2	9.9	7.0
HHS data set	(1.0)	8.0	10.9	7.7
CPU time per estimate[1]				
SAIAN data set	0.21	1.49	2.03	1.45
HHS data set	0.30	2.43	3.28	2.34
CPU time per observation				
SAIAN data set	.001	.009	.013	.009
HHS data set	.001	.010	.014	.010
CPU time per stratum				
SAIAN data set	.66	4.73	6.47	4.61
HHS data set	.15	1.20	1.62	1.15

[1]"Estimate" here refers to one mean (for a continuous variable) or a set of proportions (for all values of a categorical variable). There were a total of 35 calculated for each package.

N.B. The SAIAN data consist of 5,584 observations and 11 strata. The HHS data consist of 8,310 observations and 71 strata.

most time (71 seconds and 115 seconds). The other two packages were comparable to each other. The disparity between the two data files seems to be explained more by the number of observations than the number of strata and PSUs. Thus, the increase in the number of strata did not affect the execution time to the same extent as the increase in the number of observations.

On the PC, a different story emerges (Table 4). PC SUDAAN, regardless of data file type, was significantly faster than PC CARP. Again, it must be stated that approximate run-time was recorded, from first keystroke to final result, since PC CARP is an interactive program and response and keying speed affect the overall time. PC CARP took more than seven times longer to execute (59 minutes for SAIAN data, 126 minutes for HHS) than did SUDAAN (8 minutes and 17 minutes), primarily due to the interactive nature of PC CARP. As with the mainframe packages, the disparity in time between the two data sets is more dependent on the number of observations than the number of strata.

It should be noted that typing time required for the PC SUDAAN batch program and auxiliary files were not included in run-time, while typing in responses to PC CARP prompts was included. Both programs required substantial preparatory time with respect to syntax and content inquiries. However, once prepared to run the analyses and supply the requested information, there was a great difference in time consumed in actually executing the program. Had PC CARP allowed for batch entry of program specifications, the typing time for the batch program would not have counted in the execution time. As it now stands, PC CARP requires one to await each series of prompts, which is much more time-consuming for the user than just allowing a series of analyses to execute automatically.

4.4. *Computing costs*

As with CPU time, mainframe computing costs vary from system to system, and are

Table 4. Approximate execution times (in minutes) for microcomputer packages

	PC SUDAAN (using SAS data file)	PC SUDAAN (using text data file)	PC CARP
Total time			
SAIAN data set	8	8	59
HHS data set	17	17	126
Time per estimate[1]			
SAIAN data set	0.2	0.2	1.7
HHS data set	0.5	0.5	3.6
Time per observation			
SAIAN data set	0.001	0.001	0.011
HHS data set	0.002	0.002	0.015
Time per stratum			
SAIAN data set	0.7	0.7	5.4
HHS data set	0.2	0.2	1.8

[1] "Estimate" here refers to one mean (for a continuous variable) or a set of proportions (for all values of a categorical variable). There were a total of 35 calculated for each package.

N.B. The SAIAN data consist of 5,584 observations and 11 strata. The HHS data consist of 8,310 observations and 71 strata.

Table 5. Cost for mainframe packages (in US dollars)
[NOTE: On the NIH mainframe, cost is a function of CPU time, region used, I/O count, and number of tape drives used.]

	SESUDAAN	SUDAAN (using SAS data file)	SUDAAN (using text data file)	SUPER CARP
Total cost				
SAIAN data set	11.09	110.75	147.57	55.82
HHS data set	16.16	166.53	217.51	82.82
Comparison of cost to SESUDAAN (ratio)				
SAIAN data set	(1.0)	10.0	13.3	5.0
HHS data set	(1.0)	10.3	13.5	5.1
Cost per estimate[1]				
SAIAN data set	0.32	3.16	4.22	1.59
HHS data set	0.46	4.76	6.21	2.37
Cost per observation				
SAIAN data set	0.002	0.020	0.026	0.010
HHS data set	0.002	0.020	0.026	0.010
Cost per stratum				
SAIAN data set	1.01	10.07	13.42	5.07
HHS data set	0.23	2.35	3.06	1.17

[1]"Estimate" here refers to one mean (for a continuous variable) or a set of proportions (for all values of a categorical variable). There were a total of 35 calculated for each package.

N.B. The SAIAN data consist of 5,584 observations and 11 strata. The HHS data consist of 8,310 observations and 71 strata.

presented here to give a sense of the magnitudes of difference as well as rough estimates in terms of U.S. dollars. Mainframe computing costs for these runs were often substantial (Table 5). On the NIH mainframe computer, costs are a function of CPU time, region used, I/O count, and number of tape drives used. SESUDAAN was the least expensive to run, costing less than 20 USD for each data set. Number of observations seems to be accounting for the cost discrepancies within package for the two data files. Aside from the initial costs of hardware and software, there is no cost for the PC runs.

4.5. Computational accuracy

The estimates from all of the various packages were compared and resulted in exactly the same mean and proportional estimates as well as standard errors out to the available decimal places, usually at least four places after the decimal point. This is not surprising given that all of the packages evaluated use the Taylor series approximation to compute variances. The only exceptions were found in four of the means and associated standard errors computed in SUPER CARP; two in the HHS data, and two in the SAIAN data. However, even those means and standard errors agreed at an acceptable level (at least one place after the decimal). PC CARP estimates of means for these anomalous variables agreed with the SUDAAN packages, not with SUPER CARP. It should be noted that the PC handled the same desired level of precision as was acquired on the mainframe.

The design effects did not agree quite as precisely between the SUDAAN programs and PC CARP (SUPER CARP does not produce design effects). The disparities were mostly cases where the SUDAAN design effect was slightly higher than the PC CARP design effect, with relative differences ranging from .00% to 6.03% (aside from two outliers discussed below), and a median absolute difference of .0007. According to the formulae provided in the documentation, the variance estimates under simple random sampling assumptions (the denominator of the design effect) differ by a factor of $n/(n-1)$ between the packages. The biggest disparities occurred for mean, rather than proportional, estimates. For variables without any missing values, the design effects differed by the $n/(n-1)$ factor. For two of the continuous variables for which means were calculated, using the missing value indicator variable in conjunction with the subpopulation means option in PC CARP contributed to relative differences as high as 27%. While the formula for computing design effects for subpopulation means in PC CARP was not documented, it appears that the differential observed for PC CARP was a function of using the overall sample size and weighted population total in the calculation of the variance of the estimate under simple random sampling assumptions, rather than restricting the counts to those with non-missing data. The PC CARP "Two-Way Tables" option just produces one design effect, for the F-statistic (Test of Proportionality), and therefore was not comparable to SUDAAN.

4.6. *Tests of association in two-way tables*

Both PC CARP and PC SUDAAN can perform tests of association in two-way cross-tabulations of variables. PC CARP produces a statistic called the "Test of Proportionality," which is like a test of independence in a simple random sample setting. The test statistic has the approximate distribution of an F-statistic for large samples, and is based on a Chi-squared statistic, adjusted by the appropriate degrees of freedom. A design effect for the test statistic is also provided. PC SUDAAN produces a "Test of Independence" for each cross-tabulation. It produces test statistics based on the Wald statistic, distributed as a Chi-squared statistic with $(r-1)(c-1)$ degrees of freedom. For the example compared, the significance of the PC CARP F-statistic and the PC SUDAAN Chi-squared statistic both yielded p-values $<< .001$.

PC CARP prints out four tables for each cross-tabulation: one containing frequencies, one containing row proportions, one containing column proportions, and one with overall proportions, each shown with its estimated standard error. PC SUDAAN prints one table for each cross-tabulation, with a default printout of the elements described for PC CARP (the weighted frequencies, proportions, and associated standard errors), in addition to the unweighted cell size. One can also request that the design effects for each of the estimates be printed as well. The two packages yield exactly the same weighted frequencies, proportions, and associated standard errors.

4.7. *Documentation*

The relative ease with which one learns and uses a statistical computing package is a function of prior computing experience, statistical background, and the clarity and comprehensiveness of the documentation. Manuals should provide enough information so that the package can be learned without other assistance. Unlike the previous measures used to compare the software packages, evaluation of documentation is somewhat subjective and can vary between users.

For the most part, the software documentation was quite good for all of the packages being evaluated. Examples are used to some extent in all of the manuals, and are quite helpful when one is using one of the packages for the first time. All contain algorithms for the available analyses, for those interested in the technical aspects, although more of an effort should be made to describe why and when to use the different procedures and options. At the very least, references should be cited which would allow one to make an informed decision as to which procedures or options to use. This is particularly true of SESUDAAN and SUPER CARP, while the newer packages have improved documentation in this regard.

For a first-time user, two of the three mainframe packages required some assistance with respect to JCL (except for SESUDAAN, which has a SAS-like JCL, plus one extra line mentioned in the documentation); however, none of the packages, mainframe or PC, required any assistance other than the documentation for program execution (other than the missing value issue described earlier), even for a first-time user.

Error message documentation in the manuals, in addition to more explicit error messages in the programs to pinpoint problems, would be welcome improvements for all of the packages, since error messages are often difficult to interpret and use in resolving errors.

The authors found the SESUDAAN manual to be concise, but adequate. It has a SAS manual-like style, with clear instructions on how to structure the program statements. It is lacking somewhat with respect to why and when to use the various commands and options. The SUDAAN manual available at the time of the evaluation was designed for use with the PC and the VAX, not for the mainframe. Although the program commands are the same for the mainframe, a mainframe section would need to be added to the manual to give information on Job Control Language and file naming, and other details related to the interaction with the mainframe system. (Note that such a section was added to subsequent versions of the SUDAAN manual.) These latter issues caused the greatest problem when attempting to execute SUDAAN on the IBM mainframe. The existing SUDAAN manual, for its intended environment, is organized and clearly written.

The SUPER CARP manual, although clearly written, is quite dated. It is written in terms of "punched cards," which can be interpreted as lines of code. The manual presumes some prior knowledge of FORTRAN in its data formats and the way it reads data files and code. For example, it was not initially clear that one cannot have a blank line in the program. Although the authors seem to have abandoned this package temporarily to develop PC CARP, some attention could well be paid to updating the SUPER CARP manual, since it still is useful in cases where a file is too big for the PC to handle efficiently and effectively. Some information on SUPER CARP's interaction with the mainframe (e.g., Job Control Language and file naming/numbering) should be added as well. The PC CARP manual is very well-written and demonstrates how to use the package primarily through the use of examples. Screen displays are shown throughout the examples.

5. Summary

All of the packages evaluated for this paper proved relatively straightforward to use, after some practice. Most frustrating was the effort to correct problems which occurred in the interface between the mainframe computer and some of the mainframe software packages, due to the unhelpful error messages and the high computing costs even for SUDAAN jobs which did not run successfully. Since, in reality, many people will use these software packages without fully understanding the theory behind them, the documentation which accompanies the software should include more discussion of why and when to use certain options and procedures.

One must be very knowledgeable about the data set which is used for any of these packages prior to running the packages. While each package has different specification requirements, one should know the number of observations in the file, the number of categories for each of the categorical variables, the existence of empty PSUs for the subfile being used, the variable lengths, and the extent of missing data. Data preparation must take place ahead of time, with missing values converted to the appropriate format, sorting by nesting variables, as well as any other changes to variables. If one is using a text data file with any of these packages, any further changes to the data would also require importation back into SAS or a data base manager for modifications, then exportation back to a text file, not to mention the need to download files multiple times if the files or desired variables change.

In comparing the efficiency of the packages, one sees great variation in time, cost, and number of programming statements not only between the mainframe and personal computer packages, but between packages within each of the two environments. SESUDAAN is clearly the most efficient of the mainframe packages evaluated in terms of CPU time, dollars, and data preparation.

Run-time on the PC was difficult to standardize across packages due to the interactive nature of PC CARP. Approximate run-time was recorded, however, with PC SUDAAN running in one-seventh the time that it took PC CARP, even though hesitation at the prompted questions was minimized. PC CARP required substantially more "statements" than PC SUDAAN, but was in a menu-driven format, which may be appealing to some users of this type of software. When many estimates are required, the menu-driven format becomes less appropriate, and a batch format option would be a welcome addition. The data input requirements and output format make it more difficult to specify and read. Its inability to label categorical variable responses and the lack of p-values make interpreting output more tedious.

Only SUDAAN and PC SUDAAN had straightforward specifications for weighted means and proportions. For the other packages, all-ones variables had to be created. Indirect specifications of means, using ratios or "subgroup" means (with all observations comprising the "subgroup"), were required with SESUDAAN, SUPER CARP, and PC CARP. Similarly, for proportions, "subgroup" proportions were required with these three packages, with all observations comprising the "subgroup."

The SESUDAAN procedure is often used when computing variance estimates for weighted means and proportions for the NMES data. Based on the criteria discussed here, how do the other mainframe packages fare compared to the package currently being used? The other mainframe packages, SUDAAN and SUPER CARP, took at least seven times longer than SESUDAAN and were at least five times costlier (Tables 3 and 5). Except for SUDAAN (using a SAS data file), the other programs required at least five times the number of programming statements. However, one must consider that SUDAAN does have more features and capabilities than SESU-

DAAN, including significance tests such as the Chi-square test for independence (evaluated only briefly in this paper)[5]. There seems to be quite a sacrifice in efficiency in exchange for the added features SUDAAN offers. For all of the mainframe packages, differences in computing time and cost with different files seem to be primarily a function of the number of observations, not the number of PSUs.

Based on these same criteria, would it be feasible and desirable to use a PC package instead of SESUDAAN? Although CPU time increased from seconds to minutes, actual time from submission of the program until receipt of the output decreased. The strongest case for using the PC would be the cost issue. Some of the runs were quite expensive on the mainframe, while no computer processing costs were incurred on the PC, except for initial hardware and software purchasing costs. As for which PC package one would use, PC SUDAAN using a text data file required six times the number of programming statements compared to SESUDAAN, and PC CARP required 23 times the number of "statements" (responses to prompts). PC SUDAAN using a SAS data file seems to be the most efficient selection. As with the mainframe packages, the number of observations in the data file affects the observed differences in computing time.

The analyses that were run on the mainframe had to be run in several passes on the PC. The version of PC SUDAAN used is not equipped to make use of extended or expanded RAM. (The package's authors have added this capability in later versions.) PC CARP is also unable to make use of more than 640 kb of RAM, a restriction of the FORTRAN compiler. Although extended memory is currently of no use with these two programs, a memory manager allowing for simultaneous tasks to be performed might make using a PC more palatable for this type of analysis. Although the actual PC SUDAAN runs did not take much time, the downloading of the data and importing into SAS did occupy the PC for a significant period of time.

The scope of this evaluation is limited in that only univariate statistics (weighted means and proportions) were computed (aside from a cursory look at two-way tables and tests of association). Although many of the analyses carried out on the NMES data only require these types of estimates, the packages have many more capabilities than were examined here, particularly the PC packages. In addition, only one type of personal computer was used, and there are many other configurations being used among those who might use these packages. The evaluations are based on two data sets, with roughly 6,000 and 8,000 records. It appears that the limits have not yet been tested for these packages, with respect to file size. It should also be noted that SUPER CARP and PC CARP were not used as their authors intended, with respect to missing values. The use of missing value indicator variables was intended to get around their limitation of no missing data, but was in essence tricking the system. In using this method, the point estimate and standard errors adjusting for survey design complexities (Taylor series approach) for the responding subpopulation are correct. However, the user must be aware that the design effect is not the same as one would expect, since the standard error estimate based on simple random sampling assumptions uses the overall sample size and weighted population total in the calculation, rather than restricting the counts to those with non-missing

[5] Since not all packages had the Chi-square capability, it was decided that only univariate analyses would be performed in this evaluation, except for a brief look at tests of association with the PC packages.

data. In making national estimates, the missing values would most likely have been imputed. However, missing data often exist, particularly in exploratory analysis, and using their imputation procedure, PRE CARP, might not always be desirable.

Overall, it appears that using a PC for complex survey data analysis is certainly feasible, and may be desirable in many circumstances. One can look forward to future versions of these and other PC packages which will make even better use of the increasing capabilities of personal computers. Other features previously available only on PC CARP, such as a regression analysis capability, are now part of both PC packages, making them even more versatile. In particular, logistic regression programs for complex survey data, which are very expensive to run on a mainframe, have enormous potential as part of these PC packages.

6. References

Cohen, S.B., Burt, V.L., and Jones, G.K. (1986). Efficiencies in Variance Estimation for Complex Survey Data. The American Statistician, 40, 157–164.

Cohen, S.B., Xanthopoulos, J.A., and Jones, G.K. (1988). An Evaluation of Statistical Software Procedures Appropriate for the Regression Analysis of Complex Survey Data. Journal of Official Statistics, 4, 17–34.

Cohen, S., DiGaetano, R., and Waksberg, J. (1991). Sample design of the 1987 Household Survey (AHCPR Publication No. 91-0037). National Medical Expenditure Survey Methods 3, Agency for Health Care Policy and Research. Rockville, MD: Public Health Service.

Cox, B.G. and Cohen, S.B. (1985). Methodological Issues for Health Care Surveys. New York: Marcel Dekker, Inc.

Edwards, W. and Berlin, M. (1989). Questionnaires and Data Collection Methods for the Household Survey and the Survey of American Indians and Alaska Natives (DHHS Publication No. (PHS) 89-3450). National Medical Expenditure Survey Methods 2, National Center for Health Services Research and Health Care Technology Assessment, Rockville, MD: Public Health Service.

Fuller, W.A., Kennedy, W., Schnell, D., Sullivan, G., and Park, H.J. (1989). PC CARP. Statistical Laboratory, Iowa State University, Ames, IA, (software version used 1.3, 1988).

Hidiroglou, M.A., Fuller, W.A., and Hickman, R.D. (1980). SUPER CARP. Statistical Laboratory, Iowa State University, 6th edition, Ames, IA.

Plewis, I. (1989). Software Reviews: PCCARP/EVCARP. Applied Statistics, 38, 531–533.

SAS Institute, Inc. (1985). SAS User's Guide: Basics, Version 5 Edition. Cary, NC: SAS Institute, Inc., (Release 5.18).

SAS Institute, Inc. (1988). SAS Language Guide for Personal Computers, Release 6.03 Edition. Cary, NC: SAS Institute, Inc.

Shah, B.V. (1981). SESUDAAN: Standard Errors Program for Computing of Standardized Rates from Sample Survey Data. Research Triangle Institute, Research Triangle Park, NC.

Shah, B.V., LaVange, L.M., Barnwell, B.G., Killinger, J.E., and Wheeless, S.C. (1989). SUDAAN: Procedures for Descriptive Statistics User's Guide. Research Triangle Institute, Research Triangle Park, NC, (software version used 5.02, April 1990).

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

Received November 1990
Revised January 1993