A SURVEY OF NONRESPONSE IMPUTATION PROCEDURES [1]

David W. Chapman, Westat, Inc.

1. Introduction - A good sampling plan for a sample survey will include an extensive effort, whenever necessary, to obtain a usable response for each unit selected into the sample. Various aspects of the design, such as clustering and allocation of resources, are adapted to make this feasible and practical. However, in spite of such efforts there will always be some nonresponse in large-scale surveys. Furthermore, as indicated in a report on methodology for the Current Population Survey (12,p.53), there are no known unbiased or even consistent methods of imputing for nonresponse (unless special assumptions are made regarding the nature of the nonresponse.)

Rather than imputing for nonresponse at the time survey tabulations are prepared, tabulations could be presented with the amounts of nonresponse reported in a nonresponse category. This would allow users of the data to select their own method of making nonresponse imputations. However, the additional burden of having to compute nonresponse adjustments may not be worth having the choice of imputation method. Furthermore, users would have to make imputations from the tabulated data, and some of the related information available at the tabulation stage could not be used in these imputation procedures. It therefore appears to be more appropriate to make imputations at the time tabulations are prepared, thus eliminating the nonresponse category from the tables (except perhaps to allow for item nonresponse). If nonresponse adjustments are made, the level of nonresponse should always be reported when presenting survey results.

Pritzker, Ogus, and Hansen (10,p.445) indicate that, based on extensive experience, if the survey nonresponse rate is less than five percent, any plausible method of nonresponse imputation will probably provide acceptable results. However, in many sample surveys, nonresponse rates are substantially higher than five percent. Even with interviewer surveys which include several call-backs to households and telephone followup efforts, nonresponse rates will sometimes equal or exceed 20 percent.

In such cases, the method of nonresponse imputation can have a substantial effect on the values and biases of the survey estimates. Much research has been carried out in an attempt to discover imputation methods which reduce or minimize the nonresponse bias.

In this paper an attempt will be made to summarize some of the procedures which have been used. These procedures will be discussed in two sections: one dealing with item nonresponse, and the other dealing with total (questionnaire) nonresponse.

2. Imputation for Item Nonresponse - In most surveys some of the respondents refuse, neglect, or are unable to complete one or more questionnaire items even though they do complete most of the items. (For example, income is sometimes not supplied in a household survey.) Item nonresponse also arises when a response is received which, on the basis of the editing procedures, is determined to be unacceptable. For respondents with missing items, the information available from the completed items can be used to help impute responses for the incomplete items. In fact, sometimes there is redundant information in the questionnaire and a missing item can be inferred appropriately from other responses. In other instances, an approximate value for a missing item may be obtained by considering the general relationship between this item and another item. For example, for purchases of homes, closing costs might be estimated as a percentage of the price of the home.

Two procedures which have been used by the Bureau of Census (among others) to impute for missing values or presumably incorrect values are the "cold-deck" and "hot-deck" methods. These two procedures will be discussed in this section.

2.1 The Cold-Deck Procedure for Imputation[2] - Basically, the cold-deck procedure uses values from some prior distribution to substitute for missing responses. The distribution used is usually obtained from a previous survey taken from essentially the same population. For example, for the Census an appropriate distribution would be obtained from the previous Census, or a recent household survey.

To use a distribution of prior responses for imputation, the responses are classified by one variable or jointly by two or more variables that are reported. An attempt is made to define cross-categories (or cells) in such a way that responses will be relatively homogeneous within cells and heterogeneous between cells. There must be at least one response in each cell available for imputation. The responses in each cell are stored in the memory of the computer. A cold-deck distribution is prepared in advance for each cell.

For each missing item for a particular respondent to the current survey, the values of the appropriate completed items are noted to identify the relevant cell. The respondent is associated with the cell corresponding to the values of the items. A value is then selected from the responses in the cold deck included in the same cell. This value is usually selected at random or systematically.

As an example, suppose that the age of a respondent could be placed in a cell determined by sex and household relationship, and perhaps by the age of another member of the household. Then the age of one of the cold-deck respondents selected at random from the same cell would be inserted for the missing age.

2.2 The Hot-Deck Procedure for Imputation [3] - An objection to the cold-deck procedure is that it does not utilize data obtained from the current survey. The hot-deck procedure does use the current responses to substitute for missing items.

As with the cold-deck procedure, crossclassifications (or cells) are identified by one or more relevant variables. Initial values for each cell must be supplied from a cold deck to initiate the procedure. Then new responses are supplied for each cell from the new (or hot) deck as they appear in a pass through the file. The file may be arranged in order based on relevant variables before the procedure begins. A response remains in a cell until another respondent appears who has the same characteristicts (i.e., is in the same cell) and has a response for the particular item.

Whenever an item is missing for a respondent, he is first identified with the appropriate cell, based on the responses he does supply. Then the value retained in that cell is imputed to the respondent with the missing value. As an example, suppose that age, sex, race, household relationship, and level of education were used to define cells for imputing income values. A respondent whose income is not provided is placed into the appropriate cell as determined by his responses to the above items. The value of income in that cell (ie., the income of the respondent having the same characteristics and appearing most recently in the file sequence) would be taken as the missing income. In order to avoid using the same income value repeatedly, several income values could be stored in a cell and these values could be used in rotation, if necessary.

The method described above of imputing the preceding value for a missing item is better than using a random selection from all those in the sample falling into the same cell. This is because there is usually some special ordering of the respondents which indicates that an imputed value from a respondent close in the file would be better than one picked at random from the cell. Also, it is a more convenient procedure for computer processing.

There are some possible variations on the use of the hot-deck procedure. One would be to use as an imputed value one obtained from a regression of the particular item on several of the other items. A regression equation could be developed from either a hot deck or a cold deck.

Another variation would be to use a moving average of values in a cell to substitute for a missing value. This procedure would prevent extreme values from being duplicated and would therefore reduce slightly the variances of the estimates. However, if the ordering of the respondents were important, such a procedure would contain slightly more nonresponse bias.

3. Imputation for Total Questionnaire) Nonresponse - The hot-deck procedure described above for imputing missing items could also be used to impute values for an entire questionnaire to survey nonrespondents. As described by Pritzker, Ogus and Hansen (10,p.460), this was done in the 1960 Census by substituting for a nonresponding household the questionnaire responses of the previously listed responding household. This procedure amounts to doubling the weight[4] of the respondents whose records are duplicated. Such a procedure can yield somewhat larger variances of the survey estimates than would the procedure of weight adjustment discussed below. Hansen, Hurwitz, and Madow (3,pp.232-233) show that the maximum increase in variance is about 12 percent for the method of duplicating records.

In many surveys imputation for nonresponse is carried out by adjusting the weights of the respondents in some way to account for the nonrespondents. Alternate methods of making weight adjustments plus other methods of imputation for survey respondents will be discussed in the following sections.

3.1 The Use of a Single Weight Adjustment to Account for Nonresponse - The simplest type of nonresponse adjustment is to make one overall weight adjustment. This adjustment would be equal to the sum of the initial weights of all units selected into the sample divided by the sum of the weights of the respondents. Such an adjustment "weights up" the respondents to the total sample. (If all units selected have the same initial weights, this adjustment would equal the sample size divided by the number of respondents.)

The nonresponse bias associated with this procedure can be derived in a simple case. Suppose that a simple random sample of n units is selected from the N units in the population. The basic sampling weight for each unit selected is N/n (ie, the inverse of the selection probability). Let n_1 represent the number of the n sample units that respond to the survey. If one overall weight adjustment is used in this case, it would be n/n_1 since all sample units have the same basic weight. Also, assuming no other weight adjustments are used, each of the n_1 respondents would have the same final weight, $(N/n)(n/n_1)$.

The basic formula for estimating a population mean from weighted data is the following:

$$\bar{x} = \sum_{j=1}^{n_1} w_j x_j \Big/ \sum_{j=1}^{n_1} w_j \qquad (1)$$

where

n_1 = the number of respondents,

w_j = the final weight assigned to the jth respondent,

x_j = the value of the variable (item) for the jth respondent.

In this case, since the weights of the respondents are all equal, the estimated mean in equation 1 reduces to a simple unweighted mean of the n_1 respondents.

In this case the expected value of $\bar{x}$ is equal to $\bar{X}_{\cdot 1}$, the mean of the variable for all those in the population who would respond if selected for the survey.

The bias of $\bar{x}$ for this case can be written as follows:

$$\text{bias}(\bar{x}) = E(\bar{x}) - \bar{X} = \bar{X}_{.1} - [R\bar{X}_{.1} + (1-R)\bar{X}_{.2}]$$
$$= (1 - R)(\bar{X}_{.1} - \bar{X}_{.2}) \quad (2)$$

where

R = the population response rate (ie, the proportion of the N population units that would respond if selected for the survey),

$\bar{X}_{.2}$ = The mean of the variable for all those in the population who would not respond if selected.

As expected, the bias of $\bar{x}$ depends on two factors: (1) the population nonresponse rate, $1 - R$, which is a function of the data collection procedures; and (2) the difference between the population mean for respondents and the mean for nonrespondents, $\bar{X}_{.1} - \bar{X}_{.2}$.

In an attempt to reduce the nonresponse bias of this simple adjustment procedure, weighting classes are often defined based on the characteristics available for both respondents and nonrespondents. Separate nonresponse weighting adjustments are made within each weighting class. This procedure is discussed in the next section.

3.2 The Use of Weighting Classes to Make Nonresponse Adjustments - Suppose that the population is partitioned into c classes, based on the values of one or more survey items. Let $P_1, P_2, \ldots P_c$ represent the proportions of the population members contained in each of these classes. Also, let $R_1, R_2, \ldots R_c$ be the proportions of the units in these weighting classes that would respond if selected for the survey.

As in the previous case, suppose that a simple random sample of n units is selected from the N population units. Let $n_1, n_2, \ldots n_c$ be the number of sampling units falling into each of the classes. Of course, the n_i values are random variables and their sum must equal n. Also, let $n_{11}, n_{21}, n_{31}, \ldots, n_{c1}$ represent the number of survey respondents in the c classes. The basic sampling weight (ie, inverse of the selection probability) would be (N/n) for each sample unit (as in the previous case). However, the nonresponse adjustments would vary from class to class. For each respondent in the ith class this adjustment would equal (n_i/n_{i1}), which is the sum of the sampling weights of all sampling units falling into the ith cell divided by the sum of the sampling weights of all respondents falling into the ith cell.

The estimate, $\bar{x}_1$, of the mean would then be computed as

$$\bar{x}_1 = \frac{\sum_{i=1}^{c} \sum_{j=1}^{n_{i1}} (N/n)(n_i/n_{i1}) x_{ij}}{\sum_{i=1}^{c} \sum_{j=1}^{n_{i1}} (N/n)(n_i/n_{i1})} = \sum_{i=1}^{c} p_i \bar{x}_{i1}, \quad (3)$$

where

$\bar{x}_{i1}$ = the sample mean among respondents in the ith weighting class,

p_i = the proportion of the sample falling into the ith weighting class.

The expected value of $\bar{x}_1$ is the following:

$$E(\bar{x}_1) = \sum_{i=1}^{c} P_i \bar{X}_{i1} \quad (4)$$

where

$\bar{X}_{i1}$ = the mean of the variable for all those in the population contained in the ith weighting class who would respond if selected for the survey.

The bias of $\bar{x}_1$ can be written as follows:

$$\text{bias}(\bar{x}_1) = \sum_{i=1}^{c} P_i (1 - R_i)(\bar{X}_{i1} - \bar{X}_{i2}). \quad (5)$$

It is useful to compare the bias of $\bar{x}_1$ given in equation 5 to that of $\bar{x}$ given in equation 2. If for each of the c weighting classes $(\bar{X}_{i1} - \bar{X}_{i2})$ equals $\bar{X}_{.1} - \bar{X}_{.2}$, the biases of $\bar{x}$ and $\bar{x}_1$ are identical. Also, the bias of $\bar{x}_1$ is equal to that of $\bar{x}$ if all the class response rates, $R_1, R_2, \ldots, R_c$, are equal to the overall nonresponse rate, R.

However, if the $\bar{X}_{i1} - \bar{X}_{i2}$ values tend to be less (in absolute value) than $\bar{X}_{.1} - \bar{X}_{.2}$ and the response rates (ie, the R_i values) vary from class to class, the nonresponse bias will be reduced by the use of the weighting classes to make nonresponse adjustments. Therefore, the successful application of this procedure requires the identification of survey characteristics which will define weighting classes which vary both with respect to response rates and survey estimates. Furthermore, the characteristics used to define weighting classes must be available for both the respondents and nonrespondents. This requirement will, in many surveys, severely limit the choices of variables to use to define weighting classes.

There are many surveys in which the procedure discussed above is used to impute for nonrespondents. Among them are the Health Examination Survey (8,p.6) and the Current Population Survey (12,p.53). In Cycle I of the Health Examination Survey, seven age-sex weighting classes were defined within each of 42 primary sampling units (PSU's) for a total of 294 separate cells. Nearly half of the 294 nonresponse adjustments were between 1 and 1.10 and the three largest estimates were between 2.01 and 2.10. In the CPS, the PSU's are grouped together based on the population and labor-force characteristics of the strata from which the PSU's were selected. Within groups of PSU's respondents are placed in six cells based on race-residence characteristics.

In some cases the total number of members in each weighting class is known (or a good estimate is available) and used in the nonresponse adjustment. In such cases, the weights of respondents in a cell are weighted up to the "known" total. This procedure is closely related to stratification after sampling, discussed by Hansen, Hurwitz and Madow (3,p.232; 4, pp. 138-139). The bias of the estimate of the mean using this procedure is the same as that given in equation 5 for $\bar{x}_1$, assuming the population totals are known exactly. However, the variance and therefore the mean square error would be less for the procedure based on known totals.

Care must be taken in the application of these two imputation procedures. As demonstrated in Hansen, Hurwitz and Madow (4,pp.138-139), the variance of the estimated mean can be increased by weighting up to cell totals if the number of respondents in the cells is small. As a rule of thumb, a minimum of 20 respondents is used for the weighting cells in the CPS (12,p.53). Furthermore, for the CPS a maximum of 2.0 is taken for the nonresponse adjustment factor. In cases in which the adjustment exceeds 2, cells are combined to the extent necessary to reduce the adjustment to 2.0 or less.[5] In the Health Examination Survey (8,p.6), the 294 weighting cells average about 25 respondents each.

The choices of which variables to use to define weighting classes are usually based on which variables have the higher correlations with the zero-one response variable and with the characteristics for which survey estimates are made. It is assumed that variables which show a high correlation among respondents for survey characteristics to be estimated would show high correlations among survey nonrespondents. If so, then the use of such variables to define weighting classes would presumably minimize the nonresponse bias. The decisions on priorities of the use of variables to define weighting classes are largely subjective. These decisions involve choices as to which variables to collapse whenever weighting cells have to be combined to provide adequate numbers of respondents per cell.

A procedure which can be used to determine weighting classes objectively from a pool of possible weighting variables is the AID programmed procedure. One way this can be done is discussed in the next subsection.

3.3 The Use of the AID Programmed Procedure to Define Weighting Classes - The AID programmed procedure can be used to select which variables to use in weighting classes and also to specify which crossclassifications of these variables should be used to define weighting classes.[6] Using this procedure the sample would be divided sequentially into subgroups in a way to maximize the amount of variability explained in some dependent variable. The dependent variable used could be the zero-one response variable, or a survey questionnaire item. As a first step, the sample would be split in half based on the categories of a single variable. The variable selected from the pool of variables is the one which provides for the maximum amount of explained variance by a division into two groups. Next, one of these two groups is split again in such a way as to maximize the explained variance in the dependent variable. This procedure of defining new subgroups to account for the maximum amount of variance is repeated until the weighting classes become as small as is allowed in the specifications, or until it is no longer possible to explain meaningful proportions of remaining variance.

There has been very little investigation of the use of AID in this capacity. In a report prepared for NCHS by Chapman (2,pp.10-20), the use of AID was tested on data collected in the Health and Nutrition Examination Survey. The basic conclusion from this investigation was that the use of AID to define specific weighting classes for nonresponse adjustments does not appear to be feasible. The specific classes identified by AID can be very complex and would probably be rather awkward to work with in practice. Also, some of the classes contained a very small number of respondents, which can increase the variance (as discussed earlier). Finally, there is no easy way of merging an AID analysis based on the zero-one response variable as the dependent variable with that based on one or more survey items as dependent variables.

Perhaps the most useful information from the AID results is obtained by noting which independent variables are used most often in defining "optimal" splits in the sample subgroups. These independent variables would probably be most useful in defining weighting classes of the type discussed earlier in Section 3.2.

3.4 The "Raking" or "Balancing" Procedure for Nonresponse Imputation[7] - The "raking" procedure is one which allows the use of a large number of variables to define weighting classes simultaneously, without being concerned about the number of respondents in crossclassifications.

This method utilizes known marginal totals for the categories of two or more characteristics selected for weighting variables. These characteristics must, as before, be known for nonrespondents as well as respondents. First, the weights of the survey respondents are blown up to the given marginal totals for one of the variables. Next, the weights of the respondents, as adjusted in the prior step, are further adjusted to add to the given marginal totals for one of the other variables. This procedure is repeated for each of the variables used for the raking procedure. At this point, only the last variable dealt with will be sure to have desired marginal weight totals. However, the procedure can be repeated and the marginal totals converge to the desired numbers for all variables. The convergence proof is due to Ireland and S. Kullback(5).

The resulting adjustment applied to a particular respondent is the product of the adjustments made for the marginal total for each variable for each iteration. Estimation based on these weights has a justification in statistical information theory.[8]

As an example of the raking procedure, suppose that two variables are used in the adjustment process. Let the given marginal total for the ith category of the row variable be denoted as $N_{i.}$, and let the known total for the jth category of the column variable be noted as $N_{.j}$. Also, let $n_{i.}$, $n_{.j}$, and n_{ij} represent marginal sample size for the ith category of the row variable, the marginal sample size for the jth category of the column variable, and the sample size for the ijth cell. (These sample sizes can be taken to be sums of respondent weights.) Then, adjusting the row totals first, the new frequency (or sum of weights) of the respondents in the ijth cell is the following:

$$N_{ij}^{(1)} = (N_{i.} / n_{i.})n_{ij} \qquad (6)$$

Next, the cell frequency $N_{ij}^{(1)}$ is replaced by the following value:

$$N_{ij}^{(2)} = N_{ij}^{(1)} (N_{.j} / N_{.j}^{(1)}) \qquad (7)$$

where

$N_{.j}^{(1)}$ = the marginal total for the jth column after the first adjustment is made using equation 6.

Repeated iterations of this process can be made until the desired level of convergence on the marginal totals is reached.

3.5 <u>The Use of Regression in Weighting Adjustments</u> - A procedure using multiple regression in nonresponse imputation has been used by Astin and Molm (1) for a follow-up survey of college freshmen. Basically, the zero-one response variable is regressed on some set of independent variables which are available for both respondents and nonrespondents. The value of the regression equation for each respondent is the estimated response rate or probability of responding for population members with the same values of the independent variables. The nonresponse weight adjustment for each respondent is taken as the inverse of the value of the regression equation.

In the application of this technique the multiple correlation coefficient was less than .25. Since this indicates that only about six percent of the variation in the zero-one response variable was explained by the regression equation, there is some doubt regarding the use of this procedure. However, to compare this procedure with the weighting-class-type procedure, corresponding measures of the explained variance for the zero-one response variable would have to be observed. The low proportion of explained variation may be a result of the linearity assumption underlying the regression model which was used. The implications of the linearity assumptions are discussed for a simple case by Chapman (2,pp.60-61). A nonlinear regression model may lead to higher proportions of explained variation in the zero-one response variable. If this is the case, this method of nonresponse imputation may be more appropriate for nonlinear regression.

There are other ways that regression could be used in imputation. For example, each survey item could be regressed on the variables that are available for both respondents and nonrespondents. Of course, estimates of the regression coefficients would have to come from the respondent sample. Imputed values for the questionnaire items would be obtained for a nonrespondent from the regression equations.

A difficulty with this procedure would be the need for a large number of regression equations -- one for each questionnaire item. Another problem would be the limited information that is available for nonrespondents. That is, there may not be enough meaningful independent variables available for the regressions to be worthwhile. Perhaps this procedure would be more useful in the case of imputation for item nonresponse since a larger number of independent variables would be available in that situation.

3.6 <u>The Use in Imputation of the Amount of Effort Needed to Obtain Response</u> - If whether or not an individual selected for a sample survey participates in the survey is correlated to the measurements taken, then it seems plausible that the number of calls required to obtain participation would also be correlated to the measurements taken. If so, the number of calls required per respondent could be useful in nonresponse imputation.

One way that the number of calls could be used would be to make nonresponse weight adjustments among only those respondents who agreed to participate after several calls. Weighting classes would be defined among the nonrespondents and the "late cooperators". The late cooperators would receive weight adjustments computed in a way similar to those discussed in Section 3.2.

This procedure would minimize the bias if, indeed, the survey characteristics of the nonrespondents were more alike those of the late cooperators than those for all survey respondents. However, the validity of this assumption is questionable. It might hold for some surveys and not for others. With regard to the CPS, Waksberg and Pearl (14,p.232) indicate that there is no support for the hypothesis that the characteristics of the nonrespondents become more like the respondents as the number of visits required for interview increases. This statement was based on results from an intensive follow-up of CPS nonrespondents in which about 40 percent of the original nonrespondents were interviewed.

This procedure can have an undesirable effect on the variances of survey estimates. If the number of late cooperators is not considerably larger than the number of nonrespondents, the nonresponse weight adjustments could be relatively large. If so, the variances of survey estimates would be increased substantially.

In the imputation process another method of using the number of calls needed to complete the interview would be to try to project a mean response for nonrespondents. That is, for a particular survey item, a mean response would first be computed among respondents requiring only one call. The corresponding mean would also be computed among those requiring two calls, among those requiring three calls, etc. If the mean responses were plotted against the number of calls, a trend might be apparent.

This procedure was investigated by Chapman (2,pp.51-59) for data collected in the Health and Nutrition Examination Survey. In this case there were many different patterns observed for the various survey items. It was not possible to determine a general trend. Also for most items, the trend of mean response as a function of the number of calls was not evident enough to even attempt to project a mean value for the nonrespondents. Even for the few items for which the trend was apparent, the method appropriate to extrapolate to the nonrespondents is unclear. Consequently, the use of degree of persuasion in the imputation process for that survey did not appear to be feasible.

3.7 Imputation by Substitution of Additional Selections from the Population - For surveys in which it is feasible, imputations for nonrespondents are sometimes made by selecting substitute units from the population to take the place of the nonrespondents. For such cases an attempt is made to obtain a substitute with characteristics which are similar to those of the nonrespondent. This may be done by selecting an additional sampling unit at random from the same stratum or cluster as the unit which did not respond, or may involve a substitute picked on a subjective basis to appropriately "represent" the nonrespondent, such as a neighbor. When such a substitute sampling unit is obtained for the sample, the substitute unit is weighted as though it had been initially selected.

A possible difficulty with a substitution procedure of this type is that the effort put forth to obtain a response from each of the originally selected sampling units may not be as strong as it would have been if no substitution procedure were used. This is a serious problem since the only satisfactory way to deal with nonresponse is to keep it to a low level. Therefore every effort should be made to obtain usable responses from those units originally selected before substitutions are made.

Also, when substitutes are used, it is important to keep in mind that the total sample (ie., original respondents plus substitutes) is not equivalent to a probability sample of the same total size from the population. This is because of the bias introduced due to the use of substitutes in place of some of the originally selected units. Therefore, when this procedure is used, the amount of substitution involved should be reported.

If the above problems are taken into account and kept under reasonable control, then the substitution procedure is good if adequate substitutions are made. In particular, it has the advantage over the weighting-class method in that it does not involve any inflation of weights which causes some increase in the variances of the estimates. Also, it does provide more respondent data than the other procedures.

Of course, it is usually not possible to obtain a substitute for each nonrespondent. Therefore, even when substitution is used, one of the previous methods of adjusting weights must also be used to some extent.

As an example, Westat Research was a subcontractor to the Educational Testing Service to help design a sample of 448 elementary schools in which to administer achievement tests. The test design required exactly 448 schools. Therefore substitutes had to be obtained for each of the nonresponding schools.

The first level substitute for a nonresponding school was taken to be that school, if any, located in the same district, having the same grade structure, and having similar levels of enrollment, mean income of the surrounding community, and percent minority of the students. If such a substitute was not available, no other school in the same district was allowed as a substitute. For the second, third, fourth and fifth priority level substitutes, schools were selected at random from the same stratum as the nonresponding school.

As a result of a superb effort on the part of ETS personnel (Western Office), all 448 slots were presumably filled. Unfortunately, improper test administration forced three of these schools out of the respondent group, leaving a total of 445 responding schools. Weight adjustments using classes defined by strata were used to impute for the three missing schools.

FOOTNOTES

[1]I would like to express my sincere appreciation to Morris H. Hansen and Sidney A. Jaffe for their many helpful suggestions regarding the content of this paper. They also read over the first draft and made many useful comments.

[2]The discussion of the cold-deck imputation procedure given here is based on a description of the procedure given by Svein Nordbotten (9,pp. 26-27).

[3]The discussion of the hot-deck imputation procedure is based primarily on descriptions by Svein Nordbotten (9,pp.28-29) and the Bureau of the Census (13,pp.22-23).

[4]The weight of a respondent is a quantity which is used to give the respondent his appropriate representation in the calculation of the survey estimates. This weight consists of the product of (1) the inverse of the selection probability, (2) any ratio adjustments to known totals, and (3) nonresponse adjustments.

[5]This use of 2.0 as a maximum weight is based, to a large extent, on the overall CPS nonresponse rate of only 5 percent. In surveys with higher nonresponse rates, the maximum adjustment allowed is probably higher.

[6]A detailed description of the AID programmed procedure is given by Morgan and Sonquist (7).

[7]The general description of the raking procedure given here is based on a description by Rosenblatt (11, especially pp.4-6).

[8]This is discussed by Rosenblatt (11,p.5) and is covered in detail by Kullback (6).

REFERENCES:

1. Astin, Alexander W. and Linda D. Molm, Correcting for Nonresponse Bias in Follow-up Surveys, 1972. An unpublished manuscript available from the Office of Research of the American Council on Education, 1 Dupont Circle, Washington, D. C.
2. Chapman, David W., "An Investigation of Nonresponse Imputation Procedures for the Health and Nutrition Examination Survey," prepared for the Division of Health Examination Statistics, National Center for Health Statistics, HEW, by Westat, Inc., 1974.
3. Hansen, M.H., W.N. Hurwitz, and W.G. Madow, Sample Survey Methods and Theory, Vol. I,

New York: John Wiley and Sons, Inc., 1953.
4. Hansen, M.H., W.N. Hurwitz, and W.G. Madow, Sample Survey Methods and Theory, Vol. II, New York: John Wiley and Sons, Inc., 1953.
5. Ireland and Kullback, S., "Contingency Tables with Given Marginals," Biometrika, Vol. 55, No. 1 (March 1968), pp.179-188.
6. Kullback, S., Information Theory and Statistics, New York: John Wiley and Sons, 1959; Dover Press, 1968.
7. Morgan, James N. and John A. Sonquist, "Problems in the Analysis of Survey Data, and a Proposal", Journal of the American Statistical Association, 58 (June 1963), pp.415-435.
8. National Center for Health Statistics: Cycle I of the Health Examination Survey: Sample and Response, United States, 1960-1962. Vital and Health Statistics. PHS Pub. No. 1000 - Series 11 - No. 1 Public Health Service. Washington: U. S. Govt. Printing Office, April 1964.
9. Nordbotten, Svein, "Automatic Editing of Individual Statistical Observations", Conference on European Statisticians, Statistical Studies, No. 2, United Nations, New York, 1963.
10. Pritzker, L., J. Ogus and M.H. Hansen, "Computer Editing Methods - Some Applications and Results", Bulletin of the International Statistical Institute Preceedings of the 35th Session, Belgrade 41 (September 1965), pp. 417-441.
11. Rosenblatt, Harry M., "Study of Proposed Ratio Estimators for 1970 Census". Prepared for Census History, Statistical Research Division of the U. S. Bureau of the Census, March 11, 1971.
12. U. S. Bureau of the Census, The Current Population Survey - A Report on Methodology. Technical Paper No. 7, U. S. Government Printing Office, Washington, D. C., 1963.
13. U. S. Bureau of the Census, United States Censuses of Population and Housing, 1960: Processing the Data. Washington, D. C.,1962.
14. Waksberg, Joseph and Robert Pearl, "New Methodological Research on Labor Force Measurements". Proceedings of the American Statistical Association, Social Statistics Section, 1965, pp. 227-237.