COMPENSATING FOR WAVE NONRESPONSE IN THE 1979 ISDP RESEARCH PANEL

Graham Kalton, University of Michigan
James Lepkowski, University of Michigan
Ting-Kwong Lin, National University of Singapore

1. Introduction

The choice between weighting adjustments and imputation for handling missing survey data is generally straightforward: as a rule, weighting adjustments are used for total nonresponse and imputation is used for item nonresponses. There are, however, several situations where the choice is debatable. In general, these are situations of what might be termed partial nonresponse, where some data are collected for a sampled unit but a substantial amount of the data is missing. These situations include cases where the respondent terminates the interview prematurely, where data are not obtained for one or more members of an otherwise cooperating household (for household level analysis), and where an individual provides data for some but not all waves of a panel survey.

If weighting is used for partial nonresponse, the available responses for that unit may be employed in the determination of the weights, but the unit itself is discarded, resulting in a loss of data. On the other hand, if imputation is used, a sizeable number of responses for a partially nonresponding unit will need to be imputed, giving rise to concerns about the fabrication of much of the data and the effect of this fabrication on the relationships between variables. This paper examines the choice between weighting and imputation for handling the partial nonresponse that occurs when a respondent fails to provide data on one or more waves of a panel survey. Kalton (1985) provides further discussion of the issues involved in choosing between weighting and imputation to handle wave nonresponse, and Cox and Cohen (1985) report the results of an experimental investigation of these alternatives in the National Medical Care Expenditure Survey.

The objective of this study is to provide evidence on the choice between weighting and imputation for handling wave nonresponse in the Survey of Income and Program Participation (SIPP). The SIPP is a panel survey in which households are interviewed every four months over a period of about two-and-a-half years (Herriot and Kasprzyk, 1984). One major product of the SIPP will be an annual file combining three waves of data, and the focus of the present study is on this annual file. Since a longitudinal file for the first three waves of the first SIPP panel is not yet available, the empirical investigation reported here is based on the first three waves of the 1979 Income Survey Development Program (ISDP) Research Panel, a large-scale panel survey that was conducted as part of the development of the SIPP. All the results reported here relate only to original sample persons aged 16 and over in the area frame part of the 1979 Research Panel sample: persons sampled from the special list frames and persons joining the panel after the first wave are excluded from all the analyses.

Table 1

Person Response/Nonresponse in the First Three Waves of the 1979 ISDP Research Panel (Excluding Total Nonrespondents)

Pattern	Response (1) Nonresponse (0)	%
1	111	80.2
2	110	7.2
3	101	2.3
4	011	2.2
5	100	6.7
6	010	0.6
7	001	0.9
	Total	100.0
Number of persons		20,676

In a three-wave panel there are eight different patterns of response/nonresponse for the sampled units. Denoting 1 as response and 0 as nonresponse, one of these patterns is 000, representing the nonrespondents to all three waves. The form of adjustment for these total nonrespondents is unproblematic, namely a weighting adjustment, and hence they will not be considered further here. The distribution for the other seven patterns for the 1979 Research Panel is given in Table 1.

The first pattern in Table 1 represents those who responded on all three waves of the panel, whereas the other six patterns represent those who failed to respond on one or two of the waves. The issue under study is whether weighting or imputation should be used to handle each of these six patterns. The next section of the paper discusses how weighting adjustments might be applied, and the following one discusses the use of imputation. The final section presents some concluding remarks.

2. Weighting Adjustments for Wave Nonresponse

The use of weighting adjustments for partial nonresponse presents two additional complications beyond those that apply with weighting adjustments for total nonresponse. One results from the fact that there is a great deal more information available about partial nonrespondents than about total nonrespondents. Often only a limited amount of auxiliary information is available for total nonrespondents (such as the PSUs and strata in which they are located), whereas for partial nonrespondents there is also the information provided by their responses to the questions they have answered. The complication raised by these extra data is how they should be taken into account in determining the weighting adjustments for partial nonrespondents.

The second complication arises from the fact that surveys are subject to many different forms of analyses. Some partial nonrespondents will have provided all the data needed for certain analyses, and hence can be included in them, but they will not have provided all the data needed for some other analyses. If all those providing the requisite data for a particular analysis are included in that analysis, different analyses will be based on different subsets of the sample. This raises the complication that different sets of weights are needed according to what subset of the sample is included in a particular analysis. These two complications are discussed in turn subsequently in relation to handling wave nonresponse by weighting adjustments.

As an illustration of the first complication, consider the simple case of compensating for the second wave nonrespondents in the 1979 Research Panel. The auxiliary variables available for these partial nonrespondents are the design variables (PSUs and strata, etc.) and their wave 1 responses. The aim is to discover which, if any, of these variables are associated with response status at wave 2, and then to develop weights to compensate for differential wave 2 response rates in different parts of the sample. With the large number of wave 1 response variables, the first step in the analysis is to reduce those to be investigated in detail to a manageable number. This was done by examining the bivariate associations of each of the auxiliary variables in turn with the wave 2 response status variable. All but a few of the auxiliary variables were found to have virtually no association with wave 2 response status, and these variables were therefore excluded from the further analyses.

The next step was to employ the remaining auxiliary variables as joint predictors of wave 2 response status using SEARCH analyses (Sonquist, Baker, and Morgan, 1973) and logistic regressions. Figure 1 presents the results of a SEARCH analysis, one which explains 2.3 per cent of the variation in the wave 2 response status variable. Examination of this tree diagram shows that 88 per cent of the sample falls in cells with response rates between 87 and 92 per cent, and that 98 per cent falls in cells with response rates between 83 and 92 per cent. Only three small cells have distinctly lower response rates. In terms of weighting adjustments, giving the cell with the 92 per cent response rate a weight of 1, the weights for 88 per cent of the sample would be between 1 and 1.06 and for 98 per cent would be between 1 and 1.11. The use of these weights, with their slight variation, would be unlikely to have any appreciable effects on analyses of the data.

Figure 1

Search analysis for wave 2 response status

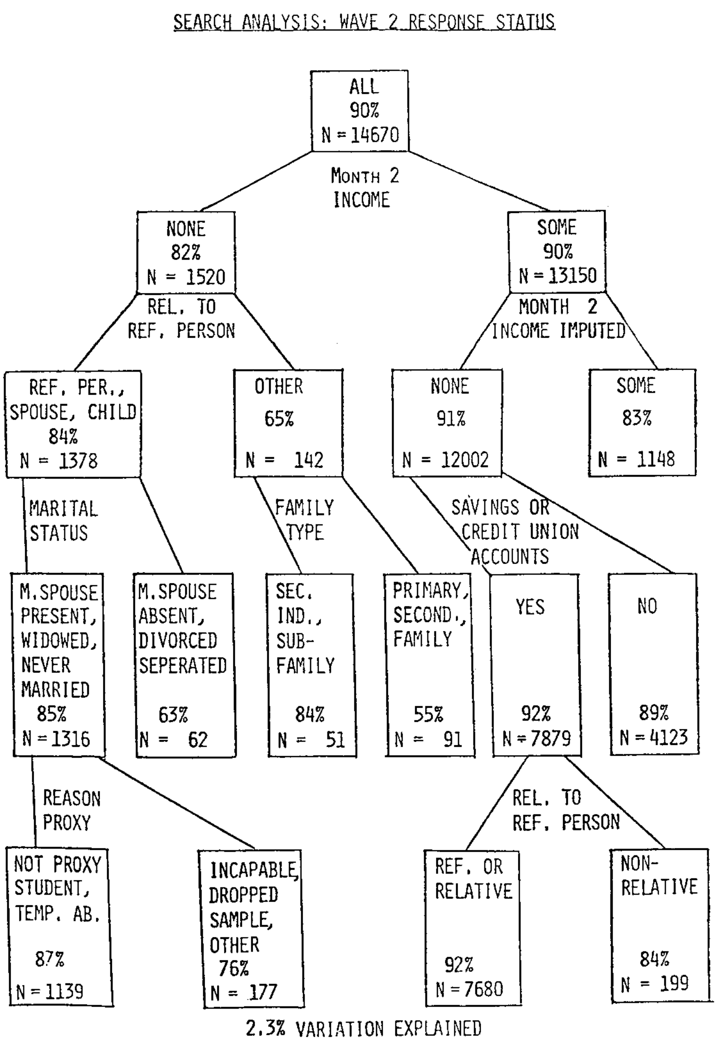

As an alternative to the SEARCH analysis, logistic regression analyses with wave 2 response status as the dependent variable were also conducted. For one of these regressions, the independent variables from wave 1 were the reason for proxy interview (1), the recipiency of interest income (2), the amount of personal earnings in month 2 (3), the relationship to the reference person (4), the type of family (5), marital status (6), and the two-factor interactions (1,2), (1,3), (1,4), (1,6), (4,5) and (5,6). Following Little and David (1983), the weights for wave 2 respondents were then set to be the inverses of their individual predicted means from this regression. Figure 2 shows the resulting distribution of weights. This distribution has a similar spread to that obtained from the SEARCH analysis, but in this case there are a few outliers. In practice, these outliers would probably be trimmed back to avoid the increase in sampling error associated with relatively large weights.

Figure 2

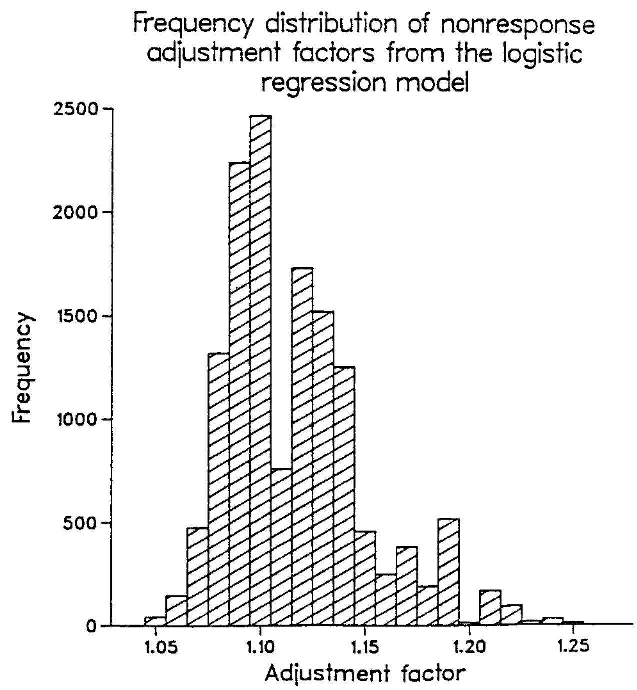

The results of the above analyses are fairly reassuring about the nature of wave 2 nonresponse. Comparisons of wave 2 respondents and nonrespondents show that the two groups are generally very similar in terms of their wave 1 responses. The differences that have been identified are not major ones, and weighting adjustments can be employed to compensate for them. Since the variation in these weights is not great, their use will not result in much loss of precision in the survey estimates. The weights from the SEARCH analysis, for example, would be likely to lead to an increase of less than 1/2 per cent in the variance of the survey estimates.

The second complication noted above concerns the need to employ different sets of weights for different types of analyses in the presence of partial nonresponse. For instance, considering the patterns of wave nonresponse in Table 1, it can be seen that patterns 1, 2, 3 and 5 provide data for cross-sectional analyses of wave 1, patterns 1, 2, 4 and 6 provide data for cross-sectional analyses of wave 2, patterns 1 and 2 provide data for analyses of changes between waves 1 and 2, and only pattern 1 provides data for forming aggregates across all three waves (*e.g.*, income over the period). For any particular analysis, the respondents in the patterns that provide the requisite data need to be weighted up to represent the other patterns. There are potentially seven combinations of waves that could be used for different forms of analysis, thus implying the need for seven different sets of weights. With more waves in the panel, the potential number of sets of weights increases rapidly. For instance, with the eight waves from a full SIPP panel, there are 255 possible combinations of waves, and hence as many as 255 different sets of weights could be required.

The number of sets of weights needed would be reduced if not all the patterns of response/nonresponse occurred. In many panel surveys the major type of nonresponse is attrition nonresponse, which refers to the situation in which a unit drops out on one wave and remains out of the panel for all subsequent waves. If the only form of nonresponse was attrition nonresponse, there would be just four response/ nonresponse patterns for a three wave panel, namely 111, 110, 100 and 000, and only three sets of weights would be needed. There would be one set of weights for each wave: these weights would apply straightforwardly for cross-sectional analyses of data from single waves, and an analysis incorporating data from two or more waves would use the weights applicable to the latest wave involved in that analysis.

Little and David (1983) propose a method for developing weights to compensate for attrition nonresponse that attempts to take account of all the auxiliary data available on the nonrespondents. The only information known about nonrespondents at the first wave (*i.e.*, the total nonrespondents) is their values on the design variables (*e.g.*, PSUs and strata), z; the information available for those who drop out at the second wave comprises their z-values and their responses at the first wave, x_1; the information available for those who drop out at the third wave comprises their z- and x_1-values and their responses on the second wave, x_2; and so on. Little and David propose running the following series of logistic or probit regressions with the response indicators r_i ($r_i = 1$ for a respondent, $r_i = 0$ for a nonrespondent at wave i) as the dependent variables:

(1) Regress r_1 on z_1 for the total sample
(2) Regress r_2 on z_1 and x_1 for respondents at wave 1
(3) Regress r_3 on z_1, x_1 and x_2 for respondents at wave 2; and so on.

The inverses of the predicted means from these regressions then give the weights needed to compensate from one wave to the next. Let these weights be denoted by w_1, $w_{2.1}$, and $w_{3.12}$. The overall weights for first wave respondents are then w_1; for second wave respondents they are $w_2 = w_1 w_{2.1}$; for third wave respondents they are $w_3 = w_2 w_{3.12}$; and so on.

Little and David (1983) also describe a weighting scheme for nonattrition nonresponse, but the simplicity of the above procedure is lost, and their scheme also has some unattractive features. As can be seen from Table 1, there were in fact a fair number of nonattrition nonrespondents in the 1979 Research Panel: the patterns 101, 011 and 001 account for 6.0 per cent of the total sample and comprise almost one-third of the partial nonrespondents. An approach that can be used to avoid the complications of the nonattrition nonresponse patterns is to convert them into attrition patterns. This can be done either by discarding some waves of data, by imputing some waves of data, or by a combination of these procedures. Thus, for instance, one might impute for the missing wave in the 011 pattern, discard the data in the 001 pattern, and either impute for the middle wave or discard the last wave in the 101 pattern. Note that if discarding is the chosen solution, the data need not have been collected in the first place (except for its potential use for methodological checks).

3. Imputing for Wave Nonresponse

When wave nonresponse is handled by imputation, all the missing items for a wave nonrespondent are assigned values, making use of responses on other waves in doing so. As Kalton and Kasprzyk (1982) discuss, the value imputed for the *ith* nonrespondent on variable y may in general be expressed as $y_i = f(x_{1i}, x_{2i},...,x_{pi}) + e_i$, where $f(\mathbf{x})$ is a function of the p auxiliary variables used in the imputation, and e_i is an estimated residual. If the e_i are set equal to zero, the imputation scheme assigns the predicted means, and the scheme may be termed a deterministic one. On the other hand, if the e_i are estimated residuals, the scheme may be termed a stochastic one. Deterministic imputations distort the shape of the distribution of y, and attenuate its variance. For this reason, stochastic imputation schemes are generally preferred.

In the SIPP and the 1979 ISDP Research Panel, in common with most panel surveys, many of the same items are repeated on each wave. Often the responses to a repeated item are highly consistent over time, and when this occurs the response on one wave can serve as a powerful auxiliary variable to use for imputing the missing response on another wave. To illustrate this point, we consider first some categorical variables and then some continuous variables from the 1979 Research Panel.

For the categorical variables we examine the consistency of responses across the first two waves of the 1979 Research Panel. The upper part of Table 2 presents unweighted cross-wave distributions of responses to whether the person worked in the quarter and to two recipiency items for original sample persons aged 16 and over who responded on both waves. The lower part of the table gives corresponding distributions of reasons for not working for those who were not at work on both waves. As the first row of the table shows, 58.2 per cent of persons reported that they worked on both waves and 34.5 per cent reported that they did not work on either wave. Thus, a total of 92.8 per cent of the respondents were consistent in their responses across the first two waves of the panel.

Table 2

Distribution of sample persons across Waves 1 and 2 for selected variables for original sample respondents for both waves ages 16 and older from the area frame, 1979 ISDP Research Panel

Item / 1st wave / 2nd wave	Yes Yes	Yes No	No Yes	No No	Consistency	Sample size
Worked in quarter	58.2	3.5	3.8	34.5	92.8	13,119
Receiving Soc. Sec.	18.4	0.4	0.9	80.3	98.7	13,151
Receiving Fed. SSI	3.2	0.3	0.3	96.2	99.5	13,151
Reasons for not working:						
Going to school	11.0	0.9	0.7	87.4	98.4	4,520
Didn't want to work	4.9	6.5	8.5	80.1	84.9	4,520
Retired	15.3	5.0	6.5	73.2	88.5	4,520

The degree of consistency of response for all the items in Table 2 is high, with the lowest level of consistency being 84.9 per cent for the responses to the item "Didn't want to work" as a reason for not working. That the "Didn't want to work" item exhibits the lowest level of consistency is perhaps not unexpected, given its greater degree of subjectivity than the other items. It is likely that all these consistency measures are underestimates, because of measurement errors, possible mismatches of respondents across waves, and other reasons. Even items like race and marital status show some degree of inconsistency. The former item has a consistency measure of 99.6 per cent, and the latter item has one of 97.8 per cent; several of the inconsistencies in marital status were in fact logical impossibilities, such as married, widowed or divorced at wave 1 and never married at wave 2.

The high levels of consistency found in Table 2 suggest that the response to one of these items on one wave is a good predictor for a missing response on the other wave. In order to illustrate how the quality of imputations based on responses to the same item on another wave may be assessed, consider the item in the first row of the table, whether the respondent worked in the quarter or not.

Among the respondents to both waves, 94.4 per cent of those who answered "Yes" to this item at wave 1 (*i.e.*, said they worked in the quarter) also said "Yes" at wave 2, and 90.1 per cent of those who answered "No" at wave 1 also answered "No" at wave 2. There were 1518 persons who answered this question on wave 1, but failed to answer it on wave 2; of these, 922 answered "Yes" at wave 1 and 596 answered "No". Using a deterministic imputation scheme, all those answering "Yes" at wave 1 would be assigned "Yes" answers at wave 2 (this being the modal wave 2 response amongst those answering "Yes" at wave 1); similarly, all those answering "No" at wave 1 would be assigned "No" answers at wave 2. Assuming that nonrespondents at wave 2 are missing at random conditional on their wave 1 responses, one can expect that 94.4 per cent of the 922 responding "Yes" at wave 1 will be correctly assigned "Yes" at wave 2 (*i.e.*, an expected 870 persons) and 90.1 percent of the 596 answering "No" at wave 1 will be correctly assigned "No" answers at wave 2 (*i.e.*, an expected 537 persons). Thus this imputation scheme may be expected to correctly assign the responses of 92.7 per cent of the wave 2 nonrespondents. Without using the wave 1 responses in the imputation scheme, all the 1518 wave 2 nonrespondents would be assigned "Yes" responses with a deterministic imputation scheme, since "Yes" is the modal answer among wave 2 respondents. Again assuming wave 2 nonrespondents are missing at random conditional on their wave 1 responses, an expected 61.2 per cent of them would be correctly assigned "Yes" responses for wave 2.

The above deterministic scheme based on wave 1 responses suffers the disadvantage that it imputes only 60.7 per cent of "Yes" wave 2 responses, whereas 61.2 per cent of "Yes" responses should be imputed to generate the correct distribution of "Yes" and "No" answers under the missing data model adopted. (The difference here is small, but it could be greater in other cases.) In addition, the deterministic imputation scheme leads to a greater stability of responses over the two waves than is implied by the model: there are no changes in responses from wave 1 to wave 2 for those with imputed wave 2 responses.

A stochastic imputation scheme can avoid these disadvantages. A stochastic scheme for the above example would assign "Yes" responses to 94.4 per cent of wave 2 nonrespondents who answered "Yes" at wave 1 and "No" responses to the other 5.6 per cent, and it would assign "No" answers to 90.1 per cent of wave 2 nonrespondents who answered "No" at wave 1 and "Yes" answers to the other 9.9 per cent. A disadvantage of the stochastic scheme, however, is that it reduces the quality of the imputations: based on the missing at random conditional on wave 1 response model, the expected percentage of correct imputations with this scheme is only 86.6 per cent.

It should be emphasized that all the measures of the quality of the imputations are based on a model for the nonrespondents. The measures may be misleading if the model fails to hold. The model used here assumes that the wave 2 nonrespondents have the same distribution of wave 2 responses as the wave 2 respondents, conditional on their wave 1 responses. Thus, for instance, it is estimated that 94.4 per cent of the wave 2 nonrespondents who answered "Yes" at wave 1 would answer "Yes" at wave 2. This estimate may be seriously in error if the model is inappropriate, and if so, the measures of imputation quality will be invalid.

Consider now the imputation of continuous variables across waves of a panel survey. Kalton and Lepkowski (1983) describe a variety of procedures that can be employed for crosswave imputation in a two-wave panel, using the value of a variable on one wave to impute the missing value of the same variable on another wave. The widely used hot-deck imputation procedure does not work well when the auxiliary variable and the variable to be imputed are very highly correlated, as will often be the case with crosswave imputation. With the hot-deck procedure, the auxiliary variable is categorized into cells, and an individual with a missing value on the variable under consideration is assigned the value of a respondent from the same cell. Thus an individual from one end of a cell may be assigned the value from a respondent at the other end of that cell. Closer matches between nonrespondents and donors can be obtained by increasing the number of hot-deck cells, but the number of cells has to be limited to ensure that matches can be made.

The categorization with the hot-deck procedure can be avoided by using some form of regression imputation. Consider, for example, the imputation of the hourly rate of pay of individual i on wave 2 (y_i) given the individual's hourly rate of pay on wave 1 (x_i). A simple regression imputation model is $y_i = a + bx_i + e_i$, where e_i is a residual term. The e_i's do not need to have a zero mean, and no restriction need be placed on their distribution. Regression imputation can be viewed as constructing a new variable $\hat{y}_i = a + bx_i$ for all individuals, imputing the e_i's for the nonrespondents, and then calculating y_i as $\hat{y}_i + e_i$. The e_i's may be assigned by any appropriate imputation scheme. They may, for instance, be imputed by a hot-deck procedure, selecting respondents' e_i's within imputation cells formed by, say, age, sex, and categorized wave 1 hourly rate of pay to assign to the nonrespondents. The choice of regression imputation model is not critical, since the assignment of the e_i's can protect against a misspecified model. The better the choice of model, however, the smaller is the variance of the e_i's, and hence the better is the quality of the imputed y_i's.

Obvious choices for a and b are the least squares estimates obtained from a regression of respondents on both waves, but simpler alternatives may also work well. The simplest model is to take $a = 0$, $b = 1$, which specifies the wave 2 value as the wave 1 value plus the change between waves: the imputation is then made for changes. Other relatively simple models set either $a = 0$ or $b = 0$; the first is a proportionate change model and the second an additive change model. There is in fact no need to include the a term in the model, since it can be incorporated as part of the residual (*i.e.*, the residual is taken to be $a + e_i$).

The quality of crosswave imputations depends on (1) the correlation between the values of the item from one wave to the next and (2) the quality of the imputations for the residuals obtained by using other auxiliary variables. We present some findings from the 1979 Research Panel relating to the first of these factors.

First consider the hourly rate of pay variable. For original sample respondents aged 16 and older in the area frame reporting hourly rate of pay on each of the first two waves of the Panel, the correlation between the two waves is 0.976. Similarly, from waves 2 to 3 the correlation is 0.964 and from waves 1 to 3 it is 0.965. (All these correlations are computed after 28 cases of apparent keying errors had been removed.)

These high correlations suggest that if a person's hourly rate of pay is available for one wave but not for a neighboring wave, the missing rate can be imputed with little error (even before considering the use of auxiliary variables in the imputation of the residual term).

Unlike hourly rate of pay, most of the amounts items in the 1979 Research Panel were reported on a monthly basis, so that there are three amounts reported for each wave. The cross-month correlations for one amount item, wage and salary income, for the first three waves of the 1979 Research Panel are given in Table 3. The data are again limited to original sample persons aged 16 and older from the area sample, and only persons reporting that they received wage and salary income are included in the correlation estimates. The correlations were computed using a pairwise missing data deletion algorithm so that the numbers of records used for different correlations may vary. Several records in the data file had apparent keying errors for the wage and salary amount (*e.g.*, the amount increased from one month to the next exactly by a factor of 10 or 100, suggesting a decimal place shift in the keying process). Since these potential errors substantially reduced cross-month correlations, the data values in error were excluded from the pairwise correlations.

Table 3

Cross-month correlations for wage and salary income amount for original sample persons ages 16 and older from the area frame, 1979 ISDP Research Panel

	1	2	3	4	5	6	7	8
2	0.903							
3	0.878	0.894						
4	0.840	0.858	0.834					
5	0.839	0.854	0.833	0.955				
6	0.828	0.853	0.816	0.945	0.944			
7	0.800	0.804	0.802	0.832	0.843	0.849		
8	0.809	0.797	0.784	0.826	0.843	0.822	0.952	
9	0.795	0.809	0.787	0.825	0.828	0.835	0.949	0.949

The correlations across months are generally high, ranging from 0.784 to 0.955. The highest correlations are between months within waves, while the lowest tend to occur for months that are more than 6 months apart. Looking down the main diagonal of the lower triangular matrix in Table 3, it can be seen that correlations between adjacent months in different waves are lower than those between adjacent months in the same wave. There are several possible explanations. One is that respondents tend to give falsely consistent responses within a wave, leading to unduly high within wave correlations. It seems more likely, however, that it is the between wave correlations that are too low. This could arise because of response variation between waves, including cases of proxy reports on one wave and self-reports on another. Also, a close examination of the records suggests that there may be some mismatched records in the file, giving rise to large differences in wage and salary income between waves.

Correlations for other amounts items in the 1979 Research Panel demonstrate similar high cross-month correlations. The correlations for wage and salary income and six other amounts items are summarized in Table 4. Average correlations were computed for the same difference between months, and separately for reports within the same wave and between different waves. For example, the average within wave correlation for a one month difference for the wage and salary amount is the average of months 1 and 2, months 2 and 3, months 4 and 5, months 5 and 6, months 7 and 8, and months 8 and 9 correlations from Table 3. The corresponding average between wave correlation is the average of the months 3 and 4 and months 6 and 7 correlations.

As observed for wage and salary income amounts, the average correlations between months in different waves for the other items are always smaller than those between months in the same wave. The correlations also decrease as the number of months between reports increases. But generally the correlations for these income items are high, indicating the kind of stability that may be used to provide accurate imputed values for missing data by using cross-month and cross-wave imputation strategies.

Table 4

Average cross-month correlations for seven amount items for original sample persons ages 16 and older from the area sample, 1979 ISDP Research Panel[1]

	One month difference			Two month difference								
	Within wave	Between wave	Within and Between	Within wave	Between wave	Within and Between	Three month	Four month	Five month	Six month	Seven month	Eight month
Wage and salary amount	0.933	0.842	0.910	0.890	0.839	0.861	0.837	0.830	0.810	0.794	0.809	0.795
Personal earnings	0.910	0.760	0.872	0.900	0.753	0.816	0.741	0.724	0.699	0.672	0.675	0.661
Social Security	0.983	0.921	0.968	0.978	0.924	0.946	0.919	0.913	0.902	0.890	0.892	0.900
Federal SSI	0.931	0.886	0.919	0.912	0.856	0.880	0.829	0.812	0.810	0.762	0.717	0.596
AFDC	0.961	0.897	0.945	0.931	0.887	0.906	0.859	0.831	0.799	0.572	0.693	0.715
Unemployment compensation	0.651	0.408	0.590	0.645	0.448	0.532	0.428	0.527	0.436	0.760	0.745	0.695
Food stamps	0.966	0.900	0.949	0.937	0.892	0.911	0.883	0.867	0.849	0.820	0.814	0.790

[1]Excluding apparent keying errors as missing data.

One of the items in the table has appreciably lower correlations than the rest, namely unemployment compensation amounts. The correlations for this item start by falling as the number of months between reports increases, but then rise for longer intervals: the correlations for months six or more months apart are in fact higher than the correlation for one month apart. This pattern of correlations may indicate that short-term unemployment receives unstable compensation while longer-term employment receives relatively stable amounts of compensation. In any case, the lower correlations for this item indicates the need for greater efforts to employ effective auxiliary variables in imputing for the residuals for unemployment compensation.

The preceding discussion has been in terms of two waves of data, one of which is missing. In a three-wave panel, the wave nonresponse patterns are 110, 101, 011, 100, 010 and 001. With pattern 110, the missing third wave data could be forecast from the second wave by one of the procedures discussed; it would probably be satisfactory to ignore the first wave data, since they are unlikely to add much explanatory power to that given by the second wave data alone. In the same way, with 011, the first wave data could be backcast from the second wave data. The missing first and third waves of data in the pattern 010 could be backcast and forecast respectively. The second wave's data in 100 and 001 could similarly be forecast and backcast, but the other missing waves are two waves apart: these could equally be imputed by one of the preceding procedures, but probably less well. The final pattern, 101, has the missing wave surrounded by nonmissing waves. In this case, it should be possible to develop a stronger imputation method, using both adjacent waves' data in the imputation scheme.

The imputation schemes described above use the response for a variable on one wave in imputing for a missing response to that variable on another wave. These schemes are especially effective when the variable is highly stable, or at least the values are highly correlated between waves, for then the observed value on one wave is a powerful predictor of the missing value on the other. A limitation to these schemes is that the value of the same variable on another wave must be available. Kalton and Lepkowski (1983) found that in many cases these schemes could not be used in imputing for hourly rate of pay in the 1979 Research Panel because a person with a missing hourly rate of pay on one wave also had a missing rate on the other wave, or was a non-wage earner or not part of the panel on the other wave. An alternative back-up imputation procedure is needed to deal with such cases, adding to the complexity of the imputations and lowering their overall quality.

Another situation giving rise to responses to the item being unavailable on another wave is when the item was included on the questionnaire for only one wave. The so-called "topical modules" on the SIPP questionnaires fall into this category. When crosswave imputation based on the same item on another wave cannot be applied, other forms of crosswave imputation, using other variables, may be employed. However, the quality of the resultant imputations will rarely compare with that of crosswave imputations based on the same item.

If imputation is used to handle wave nonresponse, the possibility of collecting data on additional auxiliary variables to improve the predictive power of the imputation models is worth considering. In particular, if a unit is a nonrespondent on one wave, additional data may be collected at the next wave. Such a strategy is being adopted in the SIPP, with the addition of a "Missing Wave" section to the questionnaire for the fourth and subsequent waves of data collection (Bailey, Chapman and Kasprzyk, 1985). This section collects information on labor force participation, income sources and asset ownership/ nonownership of respondents who, although eligible, did not respond to the preceding wave.

4. Concluding Remarks

The choice between weighting adjustments and imputation for handling wave nonresponse is not a simple one. Each method has its advantages and disadvantages. Imputation creates a completed data set that is easy for the analyst to use and, when based on a model with high predictive power, imputation is more efficient than weighting. The development of good imputations for all the variables in a missing wave is, however, a major undertaking. Unless the overall imputation scheme is constructed with great care, taking account of the cross-sectional and longitudinal interrelationships between all the variables, inconsistent or otherwise unacceptable imputed values may be assigned. In any event, imputation fabricates data to some extent and it will cause an attenuation in some of the covariances between variables. The amount of fabrication and attenuation is slight when powerful crosswave imputation models are used, but such models cannot be used in all cases. On the other hand, while weighting avoids the attenuation problem, the need to use different sets of weights for different types of analyses creates complexities for the analyst and can lead to inconsistent results. With both imputation and weighting having their advantages and disadvantages, it may be that some combination of the two methods, such as that outlined at the end of Section 2, is the best solution.

References

Bailey, L., Chapman, D.W. and Kasprzyk, D. (1985). Nonresponse adjustment procedures at the Census Bureau: a review. *Proceedings of the Bureau of the Census First Annual Research Conference*, 421–444.

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

Herriot, R.A. and Kasprzyk, D. (1984). The Survey of Income and Program Participation. *Proceedings of the Social Statistics Section, American Statistical Association*, 107–116.

Kalton, G. (1985). Handling wave nonresponse in longitudinal surveys. *Proceedings of the Bureau of the Census First Annual Research Conference*, 453–461.

Kalton, G. and Kasprzyk, D. (1982). Imputing for missing survey responses. *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 22–31.

Kalton, G. and Lepkowski, J. (1985). Following rules in the Survey of Income and Program Participation. *Journal of Economic and Social Measurement*, 1, (forthcoming).

Kalton, G. and Lepkowski, J. (1983). Cross-wave item imputation. In *Technical, Conceptual, and Administrative Lessons of the Income Survey Development Program (ISDP)*, M. David, ed., pp. 171–198. New York: Social Science Research Council.

Kasprzyk, D. and Kalton, G. (1983). Longitudinal weighting in the Income Survey Development Program. In *Technical, Conceptual, and Administrative Lessons of the Income Survey Development Program (ISDP)*, M. David, ed., pp. 155–170. New York: Social Science Research Council.

Little, R. and David, M. (1983). Weighting adjustments for non-response in panel surveys. Bureau of the Census working paper.

Sonquist, J.A., Baker, E.L., and Morgan, J.N. (1973). *Searching for Structure*. Ann Arbor, Michigan: Institute for Social Research.