METHODOLOGICAL ISSUES IN THE ESTIMATION OF THE DISTRIBUTION OF HOUSEHOLD NET WORTH: RESULTS FROM THE 1989 SURVEY OF CONSUMER FINANCES

Arthur B. Kennickell, Board of Governors of the Federal Reserve System
R. Louise Woodburn, Statistics of Income Division, Internal Revenue Service
Mail Stop 180, Washington, DC 20551

Key Words: Complex Surveys, Weight and Variance Estimation, Unit and Item Nonresponse

In the U.S., we have very little representative data on the wealth of individuals. There are three different approaches that have been used to estimate wealth: directly using survey data, capitalizing income flows from survey or administrative data, and extrapolating wealth data reported for decedents on federal estate tax returns. Each of these approaches gives rise to difficult issues to resolve and various assumptions to accept in the development of the estimates. Greenwood [1983] reports on results using the income capitalization method. In using this method, assumptions must be made concerning the capitalization rates and the income flows to include. Extrapolating wealth data of decedents is referred to as the estate multiplier technique and is discussed in Johnson and Schwartz [1990]. In developing such estimates, there are many difficult issues to resolve, such as, do individuals change their portfolios and arrangement of their assets and liabilities within families in anticipation of death, how should the wealth of the decedent group be used to generalize about the wealth of the entire population. Johnson and Woodburn [1992] address these issues and detail recent improvements in the estate multiplier technique. The measurement of wealth using survey data is the topic of this paper.

There are three major problems in using surveys for the measurement of wealth. First, wealth is a construct based on the sum over a complicated array of assets and liabilities that may be difficult to explain to survey respondents. Second, wealth is relatively concentrated, and, thus, sampling to measure wealth faces many of the same problems faced in measuring rare events. Without some special provisions to target wealthier households, a random sample is likely to be a very inefficient vehicle for measuring wealth. Finally, in surveys there are systematic nonresponse problems, both at the unit and the item level. Where nonresponse rates are high, as they are in survey measurements of wealth, corrections for nonresponse are critically important.

In this paper we provide an overview of some of the interesting problems of measuring household wealth using survey data. While we address some theoretical problems, our main focus here is on the practical problems we have encountered in the 1989 Survey of Consumer Finances (SCF), a household wealth survey sponsored by the Federal Reserve Board (FRB) in cooperation with the Statistics of Income Division (SOI) of the Internal Revenue Service.[1] Further background on the survey and additional detail on specific methodological issues may be found in Heeringa and Woodburn [1992], Kennickell [1991a, 1991b], and Kennickell and Woodburn [1992a, 1992b].

The remainder of this paper is divided into four parts. First, we briefly discuss the design and execution of the SCF questionnaire. Next, we deal more extensively with the design of an appropriate sample for the measurement of wealth. Third, we address issues of estimation using the resulting data. We conclude with a brief summary and a discussion of our plans for future work.

Questionnaire Design and Data Collection

The goal of the SCF is to collect information on household assets and liabilities from a representative sample in order to address research questions involving wealth. The questionnaire for the 1989 survey was designed to provide a detailed view of the financial position of households. Respondents were asked for information on the types and amounts of assets they held. For household liabilities, they were asked the types of loans, the purposes of their borrowing, the amounts they still owed, and the detailed terms of payments of their loans. In addition, households were asked for information on the institutions where they held their assets and loans, on their current jobs and pension rights, and on many demographic characteristics such as marital history, family composition, and health status. Data for the survey were collected between the months of July 1989 and March 1990 by the Survey Research Center at the University of Michigan. Interviews were mainly conducted in person and averaged about 77 minutes, though many interviews were much longer.

Such an interview places a tremendous burden on both respondents and interviewers. Respondents must deal in great detail with a subject that many people find complicated and confusing, and interviewers must be able to understand the subject well enough to deal with respondents' questions without disrupting the interview.

Inevitably, the conceptual frameworks of researchers and survey respondents overlap imperfectly. This problem is particularly acute in the SCF, in which respondents with a very great range of financial sophistication must all be asked the same set of core questions. Thus, an important goal of the SCF questionnaire design is to minimize this "conceptual variance" through questionnaire design and interviewer training, and to eliminate as much as possible of the remaining misclassification and other misreporting (particularly double-counting) after data collection by using trained editors.

The SCF is quite long, the questions ask for much detail, and some respondents view the subject matter as intrusive. Even with great attention to questionnaire design, some respondents will be unable or unwilling to answer some questions. Every observation in the survey contains at least one piece of missing information, though often only a trivial item such as the interviewer ID number is missing. Some respondents who were reluctant to provide dollar answers were allowed to respond with letters corresponding to ranges provided on a card shown to the respondent. Excluding these range responses, the mean number of missing values per case is 21.6. The maximum possible number of missing values is about 6 million.[2]

Sample Design

Household net worth is highly skewed. Estimates prepared by Kennickell and Woodburn [1992a] based on the 1989 SCF suggest that the top one-half percent of the household net worth distribution accounted for about 29 percent of this variable. Moreover, some assets and liabilities – e.g., investment real estate loans, stocks, etc. – are highly correlated with net worth. Thus, for efficiency in the measurement of such variables, wealthier households should be heavily over-sampled. Other components of net worth – e.g., mortgage debt, automobiles, credit card debt – are much more broadly distributed in the population. Such items are better represented by something closer to a simple random sample.

To provide good coverage of both concentrated and more broadly-distributed assets and liabilities, the SCF sample is based on a unique dual frame design[3] that relies in part on the use of administrative data.[4] One part of the sample is a straightforward multi-stage area-probability cluster sample. The other part of the sample, a list sample, is more complicated and is used to systematically supplement the area sample.

In principle, we would like to use wealth information to ensure that the sample is efficiently designed for measuring wealth. Since in the U.S. there is not a general population register of wealth available, we make use of a proxy, a "wealth index" in the design of the list sample. The list sample is derived from the 1987 Individual tax file which is developed by SOI. This file was made available through an inter-agency agreement and a well-defined contract with our survey vendor.[5] The wealth index is computed from capitalized income flows observed in the 1987 Individual file. Income flows from items, such as, businesses, interest, and dividends, are capitalized using fixed rates of return intended to reflect the average rate for the assets underlying each income type. For example, if a return showed $100,000 in interest income, it was assumed that the individual held $1,000,000 in interest-yielding assets based on a %10 annual rate of return. Returns were sorted into seven strata based on the wealth index, and returns in the different strata were sampled disproportionately.

There are several non-trivial problems in designing the SCF list sample using the wealth index. First, not all assets generate income that would be recorded on an individual tax return and, in general, debts are not reported either. For example, it is only for homeowners who have mortgages and who itemize their deductions that we know anything about an individual's house. Other than mortgages on a principal or secondary residence, there is generally no systematic trace of any other borrowing. No estimate of debt is incorporated in the wealth index. Second, using a fixed rate of return may also be incorrect. Particularly in the case of closely-held businesses or newly started businesses, the current income flows may be a poor indicator of the value of the business. Rates of return also vary by risk classes of assets. While market prices of assets should change to equalize rates of return within risk classes, we are unable to observe risk characteristics. Third, particularly in the case of very complicated finances, the flows of income within only one year may be a very poor indicator of a taxpayer's longer-run income. For example, gains and losses may be clustered for tax purposes. For this reason, some use of longitudinal data in the sample design might be useful; we intend to explore this possibility with the 1995 SCF.

There are also some conceptual differences between the area and list frames. The elements of the area-probability frame are dwelling units at the time of the survey. The list frame differs in three key ways. First, the frame elements are income tax returns filed by individuals or couples. In some households, couples may file separate returns, or there may be other individuals in the primary economic unit of the household who also file returns. The area-probability frame would sample these tax-payers as a unit. However, it appears that the problem of multiple filers is not a serious one among the 1989 SCF list respondents with only 1.6% of the list sample respondent couples reporting to have filed separately for tax year 1988. Second, because by definition, a person must have filed a tax return to be eligible for selection into the list sample, that sample is not representative of nonfilers. According to the 1989 SCF, about 13% percent of households did not file any sort of individual tax return. As would be expected, these families tend to be much poorer than the typical family. Third, because the list sample is drawn from a sample of 1987 returns, it misses households that may have been created since 1987.

Thus, the list frame is a noisy approximation to the desired frame – a "fuzzy frame." However, with the exception of the exclusion of nonfilers and families formed since 1987, a sample drawn from this frame should not be biased in its representation of the general population, only an inefficient one.[6]

The achieved cross-section sample from all parts of the design includes 3,143 families, of which 866 come from the list frame.[7] The area-probability cases were approached directly by interviewers and the response rate for these cases was about 69 percent. In contrast, the list-sample cases were given a prior opportunity to refuse participation by returning a postpaid card. About 36 percent of the original sample of list cases refused participation at this stage by returning the card. The remainder of the list cases were approached by interviewers, yielding an overall interview rate for the list sample of about 34 percent. Response rates by wealth index stratum decrease from 49% for the lowest index stratum to 20% for the highest. The decreasing response rates indicate that the ultimate probability of observation differs systematically from the selection probability.

Given the high level of unit nonresponse in the list sample, one might reasonably question the representativeness of that sample. Consequently, one might also question the use of the data for estimating characteristics of the upper end of the distribution of wealth that this sample is intended to reflect. It is important to note that this is not a problem unique to the SCF. While the noninterview rate for the list cases is high according to usual criteria, this figure merely makes explicit what is latent in other household surveys without such a sample or other auxiliary information to identify the problem. Fortunately, in the SCF we are also able to use the auxiliary information in the list frame to make adjustments to the sampling weights to compensate for nonresponse.

Estimation

In this section we deal with the problems of estimation given a dataset with significant item and unit nonresponse. Ultimately, we would like to compute a broad range of estimates and to have the ability to compute standard errors in a straightforward way. There are at least two possibilities for proceeding.

We could perform model-based estimation designed specifically for each problem incorporating frame information and other information. Given the very diverse demands on the dataset and the practical limitations on the time available to devise complex estimators, this approach presents tremendous difficulties. A second possibility – the approach we take here – is to create a fully imputed dataset and compute adjustments to the sampling weights to compensate for unit nonresponse.

Nonresponse-Adjusted Sampling Weights

To assess the variability of our calculations with respect to alternative assumptions about weighting, two methodologies were used to develop the weights for the 1989 SCF, resulting in a model-based weight and a design-based weight. The design-based weight we describe here is a later generation of the weight used in the calculations reported in Kennickell and Shack-Marquez [1992] and is based on the original design probabilities developed by Heeringa and Woodburn [1992]. The model-based weight was used in calculations reported in Kennickell and Woodburn [1992a] and its construction is described there in more detail.

A weighting scheme for the 1989 SCF must combine information in the list and area samples. Following in the spirit of Hartley [1962], one dual-frame weight that pools the two samples entails knowledge of the probability of observation of each case under both designs (assuming independence), given by

$$(1)\ W^i = \{\pi^i_{AP} + \pi^i_{LIST} - \pi^i_{AP}\,\pi^i_{LIST}\}^{-1},$$

where π^i_{AP} is the probability of observation of a case within the area-probability frame and π^i_{LIST} is the probability of observation of a case within the list frame. The probability of observation for a case is the probability of selection multiplied by the probability of response given selection.

The methodologies used to develop the two types of weights share a common framework.

For both methodologies, the dual-frame weight according to equation (1) is computed and a common post-stratification procedure is used. The two approaches differ only in the determination of the probability of observation for list and area-probability sample cases under the list design.

Design-Based Probabilities of Observation

The term "design-based" is used loosely here. In fact, it is not possible to compute a pure design-based weight for two reasons. First, because it is not possible to match the area-probability cases to tax records, their exact probability of selection under the list frame is not known. Second, in no case is the probability of response known.[8]

For area and list sample respondents, respectively, the probabilities of selection are taken directly from the original sample design described in Heeringa and Woodburn [1992]. In each of the samples, nonresponse was assumed to be completely random within adjustment classes. For the area-probability sample, the adjustment classes were geographic areas, and for the list sample, the classes were the six original wealth index strata. Uniform non-response adjustments were made within each non-response class to estimate the probability of observation.[9]

Because the probabilities of observation of area-probability cases under the list design and of list cases under the area-probability design are not known with certainty, some modeling assumptions must be made. For each list case, the probability of observation under the area-probability design was assumed to be equal to the median probability of observation of the area-probability cases in the same geographic area. Given the differential nonresponse we observe in the list sample across the wealth index strata, we can be almost certain that for wealthier families in the list sample, this procedure will overstate the probability of observation.

Estimation of the probability of observation of an area-probability case within the list frame is more complicated. Using the survey data, broad wealth classes were defined for both the list and area-probability cases. Area cases in a given class were assigned the median weight of list cases in the same class, and the area weights were forced to sum to control totals in these classes developed from the list sample.

Model-Based Probabilities of Observation

For the model-based weights, the probability of observation of an area-probability case under the area design is taken to be the same as in the case of the design-based weights. We have no information to improve on this assumption.

However, in sharp contrast to the design-based weight, the calculation of the probability of observation in the list frame of the list sample cases ignores the original SCF sample design almost entirely. The list sample cases were drawn from the SOI 1987 individual tax file, which is itself a sample – albeit a very large sample – from the universe of U.S. individual tax-filers.[10] A great deal of information is known about both survey respondents and nonrespondents in the list frame.[11] Our strategy is to use this auxiliary information to estimate the representation in the population of each list observation using the entire SOI file to estimate a response propensity model.[12] Formally, we compute a probit model using the entire SOI file (not just the sample selected for the 1989 SCF) to estimate the conditional probability of observation within the SOI file, given characteristics present in that file.[13] The probability of observation for a list case under the list frame is given by the product of the conditional probability predicted by this model and the probability of selection into the SOI sample. A special adjustment was made in the calculation for observations with negative net worth. The predicted list sample weights were then forced to sum to the stratum totals in wealth index stratum that represent the 1989 population of households filing a tax return.

We are able to do more limited modeling of the implied list weights for area-probability cases. Area-probability cases that did not file a tax return as reported in the interview are given a zero list weight. For the remaining cases, those having $100,000 and more of net worth are treated separately. For these wealthier cases, we model the propensity-score-adjusted list weights directly in terms of a number of survey variables. The model was estimated in logs using the list sample cases with net worth of $100,000 and more and was used to predict weights for area-probability cases.[14] The predicted weights were adjusted to sum to a population total estimated for this group using the list sample.

Largely because the list sample is very thin at the bottom of the net worth distribution, estimates of the list weight model that included the entire list sample produced very unstable values for the predicted list weights for lower-wealth families. The remaining area sample cases with less than $100,000 of net worth were assigned the median propensity-score-adjusted list weight for list cases with less than $100,000 of net worth. The weights for these observations were adjusted to sum to the 1989 population total less the estimated number of nonfilers and the number of cases estimated to have $100,000 or more of net worth.

The probability of observation of list sample cases under the area design is assumed to be zero for list cases with $2.5 million or more of net worth and for those with negative net worth.[15] Other list cases are assigned the median area-probability weight of area-probability cases within the same region and MSA type.[16] The probabilities were rescaled so that the implied weights summed to the 1989 population total minus the estimated number of nonfilers and minus the estimated number of observations with $2.5 million or more of net worth estimated using the propensity-score-adjusted list weights.

Combined Dual-Frame Weights

The component probabilities were merged using the formula given by (1). These combined weights were subjected to post-stratification and raking.[17] First, the weights were adjusted to reproduce exactly the number of families estimated by the list sample to have net worth in the categories $1 million to $2.5 million, $2.5 million to $10 million, and $10 million to $250 million. List and area-probability observations having net worth of $1 million or more and list cases with negative net worth less than -$100,000 were excluded from further adjustments. The remaining observations were raked for three iterations to distributions of age by homeownership and region by MSA status.

While our means of evaluating the reliability of the weights is quite limited, we can make use of the list frame data to compare the distributions of key income variables estimated using the list sample weights to the base distribution computed with the entire list frame.[18]

Multiple Imputation of Missing Data

As noted above, we have dealt with the problem of item nonresponse by imputing the missing data. Missing values are multiply-imputed using Gibbs' sampling. The general imputation algorithm (FRITZ) is described in Kennickell [1991a].

Variance Estimation

The estimates computed using the 1989 SCF presented here challenge traditional variance estimation methodology. Without many simplifying assumptions, traditional variance estimators would not be appropriate given the complexity of the sample design. Also, we need to estimate variances for nonlinear and nonfunctional statistics – for example, the proportion of wealth held by the wealthiest half percent of the population. In order to compute variances for such estimates, it is necessary to use a replication technique, such as the bootstrap, balanced repeated replication, or the jackknife.[19] In order to facilitate the estimation of variances due to sampling and the uncertainty inherent in the computation of the weights, we have chosen to compute a set of replicate weights based on the bootstrap technique.[20] We drew eleven bootstrap samples for this purpose using all available cross-section interviews as a base.[21] The procedure we applied to draw each bootstrap sample treats the 1983 panel cross-section, 1989 area-probability, and 1989 list samples separately and attempts to mimic the major sources of variation for each sample.

We repeated the procedures used for the construction of the model-based weights for each bootstrap sample to yield eleven replicate weights. Thus, for each bootstrap sample, we recomputed all of the necessary models and subjected each bootstrap sample separately to the post-stratification process. While we believe that the bootstrap weights provide an adequate representation of variance due to sampling, the theoretical properties of variance estimates based on bootstrap weights for complex surveys are still being developed.[22] The principal virtue of the bootstrap weights we have computed here is that they are straightforward to compute and they are relatively simple to use in estimation.

Conclusions and Future Research

In this paper we have dealt with some of the larger methodological issues of wealth estimation using the 1989 SCF. One of the major on-going efforts in the SCF is to convince users to take the trouble to estimate variances for their estimates. To this end, we have developed multiple imputations to be used for estimating imputation variances, and bootstrap replicate weights to be used for estimating sampling variance.

The creation of a public use dataset raises the very important and non-trivial issue of confidentiality of the survey respondents. In a highly stratified survey, such as the SCF, where the oversampled group provokes strong curiosity in many people, it is important to take special precautions to protect the identity of survey respondents. While there are methods for limiting information in tabular presentation, the development of practical ways of protecting large, complicated datasets is still at a very early stage. The procedures we have adopted for the 1989 SCF are fairly crude -- suppression of critical variances, collapsing cells for categorical variables, constrained imputation of sample outliers, and other adjustments. Although some interesting partial solutions have been proposed, until we have a clear legal way of handling disclosure risks, we will remain cautious concerning the release of micro data.

Finally, an important part of the 1989 SCF sample design, which we have only mentioned here, is the 1989 panel sample based on the 1983 SCF. A part of the 1989 sample has both cross-section and panel representation, and a part has only panel representation. An important subsample that is taken to have only panel representation is the 1983 list sample. It is unknown what representation this group might have in 1989. It is possible that we may be able to develop model-based weights to estimate that representation if we can collect sufficient data on the sample. Construction of proper weights is also complicated by nonresponse in 1989 and by splitting of households.

Another major problem for dealing with the panel data is imputation. In principle, imputations should condition on all available information. When the original 1983 imputations we estimated, there was no 1989 information. After the 1989 SCF, both the 1983 and the 1989 imputations should be conditioned on a common database. While this processing raises no large problems of imputation theory, there are enormous practical problems that grow more complicated with each additional wave of a panel. We are dealing with panel imputation in the 1989 SCF by creating a very reduced version of the 1983 SCF that will be reimputed simultaneously with the 1989 data. Much clever innovation is needed to make practical progress here.

ENDNOTES

1. The 1989 SCF is one in a series of household surveys conducted by the Federal Reserve Board to collect data on household finances. Results from earlier surveys can be found in Avery, Elliehausen and Kennickell [1988], Avery and Kennickell [1992] and Projector and Weiss [1966].

2. If one looks only at dollar amounts of financial assets, out of a maximum of 136,908 data items, 3350 are missing.

3. See Heeringa and Woodburn [1992] for additional details.

4. Administrative data is often used in the aggregate to improve survey estimates. Huggins and Fay [1988] describe such an application for SIPP longitudinal estimation.

5. For a description of the Statistics of Income Individual Program in place when the SCF selections were made, see Individual Income Tax Returns, 1987 [1990]. Statistical and research uses of SOI data are closely regulated to guarantee that individuals (and other entities) will remain protected against any disclosure of their financial and tax data [e.g., Wilson and Smith, 1983]. For the 1989 SCF, contractual agreements between the Federal Reserve Board, the Survey Research Center, and SOI clearly specify the limitations on the use of the administrative data in order to guarantee the privacy rights of the individual taxpayers.

6. To evaluate the efficiency of the sample design, we computed a regression model of the computed wealth index with several survey variables. Not surprisingly, the index is a better proxy for gross assets than net worth. More details are given in Kennickell and Woodburn [1992b].

7. The 1989 SCF design included an overlapping panel/cross-section sample, based in part on the design of the 1983 SCF. This paper deals only with the cases that have cross-section representation in 1989.

8. This second point is one that holds true for most surveys. Typically, adjustments are made uniformly within adjustment classes. This adjustment assumption is a model-based one, though it is sufficiently associated with elementary sampling that many people do not consider this. Where there is limited auxiliary information, ratio adjustments may be the best that can be done.

9. Control totals for the list sample were adjusted to represent the 1989 population of households filing tax returns rather than the 1987 population of tax filers. An estimate of the number of nonfilers derived from the area frame was used to adjust the overall total.

10. The sample design for the 1987 SOI Individual study can be found in Individual Income Tax Returns, 1987 [1990].

11. Among the very few pieces of survey information that have been connected to the list frame is response status of each sample case. It should be emphasized that while we do know the set of cases that responded, we have no knowledge of which interviews correspond to which administrative records.

12. Little and David [1983] use the propensity score method to adjust for nonresponse in panel survey settings. More recently, Ekholm and Laaksonen [1991] discuss the application of propensity scores in the Finnish Household Budget Survey.

13. The estimated model is given in Kennickell and Woodburn [1992b].

14. The resulting model is given in Kennickell and Woodburn [1992b].

15. This assumption likely understates the probability of observation for higher-wealth observations since there are a small number of relatively wealthy area sample cases. While one might argue for a higher cut-of, such a change has little effect.

16. Several possible models for these probabilities were investigated. Unfortunately, the models are too sensitive to be useful for weight adjustment.

17. A description of the raking ratio estimation technique can be found in Oh and Scheuren [1978].

18. It is important to note that this does not

entail the merging of survey data with the SOI file. Only a response indicator is passed to a version of the SOI file held at the Board of Governors. Further details are presented in Woodburn [1991].

19. An overview of replication techniques for variance estimation and applications in complex surveys can be found in Skinner et. al. [1989].

20. This construction is described more fully in Kennickell and Woodburn [1992b]. The computation of total variance is based on the replicate bootstrap weights and the uncertainty due to imputation of missing values. The incorporation of the imputation variance follows the process outlined in Rubin [1987].

21. The number of bootstrap samples generally recommended is far larger than 11 (12 including the full sample weight). Applications of the bootstrap technique to complex survey settings is fairly new. Typically, such applications involve the selection of 50 to 500 bootstrap samples to compute variances and confidence intervals for a specific estimate. We have chosen to compute weights to correspond with our bootstrap samples and to provide these weights to the data users as a variance computation tool.

22. Rao and Wu [1988] compare the performance of the bootstrap with other resampling methods of computing variances in complex survey settings.

BIBLIOGRAPHY

AVERY, R. B., ELLIEHAUSEN, G. E., and KENNICKELL, A. B. (1988). "Measuring Wealth with Survey data: An Evaluation of the 1983 Survey of Consumer Finances," Review of Income and Wealth, Series 34, No. 4, pp. 339-369.

AVERY, R. B., and KENNICKELL, A. B. (1992). "Household Saving in the U.S.," Review of Income and Wealth, Series 37, Number , pp. 409-432.

EKHOLM, A. and LAAKSONEN, S. (1991). "Weighting via Response Modeling in the Finnish Household Budget Survey," Journal of Official Statistics, 7, 3, 1991, pp. 325-337.

GREENWOOD, D., (1983). "An Estimation of U.S. Family Wealth and its Distribution From Micro-Data, 1973," Review of Income and Wealth, March, pp. 23-44.

HARTLEY, H. O. (1962). "Multiple Frame Surveys," Proceedings of the Social Statistics Section, American Statistical Association, pp. 203-206.

HEERINGA, S., JUSTER, F. T., and WOODBURN, R. L. (1991). "The 1989 Survey of Consumer Finances: A Survey Design for Wealth Estimation," Statistics of Income and Related Record Research, 1990-1991, Internal Revenue Service.

HEERINGA, S., and WOODBURN, R. L. (1992). "The 1989 Survey of Consumer Finances, Sample Design Documentation," (forthcoming).

HUGGINS, V. J., and FAY, R. E., (1988). "Use of Administrative Data in SIPP Longitudinal Estimation," Proceedings of the Section on Survey Research Methods, American Statistical Association.

INTERNAL REVENUE SERVICE, (1990). Individual Income Tax Returns 1987, Department of the Treasury, pp. 13-17.

JOHNSON, B. W., and SCHWARTZ, M., (1990). "Estimates of Personal Wealth, 1986," Statistics of Income Bulletin, Spring, Volume 9, 4, pp. 63-78.

JOHNSON, B. W., and WOODBURN, R. L. (1992). "The Underlying Methodology of the Estate Multiplier Technique: Recent Improvements and Estimates for 1989," Working Paper, Internal Revenue Service.

KENNICKELL, A. B., and SHACK-MARQUEZ, J. (1992). "Changes in Family Finances from 1983 to 1989: Evidence from the Survey of Consumer Finances," Federal Reserve Bulletin, Board of Governors of the Federal Reserve System, 78, No. 1, 1992, pp. 1-18.

KENNICKELL, A. B. (1991a). "Imputation of the 1989 Survey of Consumer Finances:

Stochastic Relaxation and Multiple Imputation," Proceedings of the Section on Survey Research Methods, American Statistical Association.

KENNICKELL, A. B. (1991b). "Technical Codebook for the 1989 Survey of Consumer Finances," Board of Governors of the Federal Reserve System.

KENNICKELL, A. B., and WOODBURN, R. L. (1992a). "Development of Model-Based Weights for the 1989 Survey of Consumer Finances," Technical Documentation. (forthcoming).

KENNICKELL, A. B., and WOODBURN, R. L. (1992b). "Estimation of Household Net Worth Using Model-Based and Design-Based Weights: Evidence from the 1989 Survey of Consumer Finances," Working Paper, Board of Governors of the Federal Reserve System.

LITTLE, R. J. A. and DAVID, M. H. (1983). "Weighting Adjustments for Nonresponse in Panel Surveys," Working Paper. Washington D.C., U.S. Bureau of the Census.

MCCUBBIN, J., (1990), "The Intergenerational Wealth Study: Basic Estimates, Data, 1916-1945," Statistics of Income Bulletin, Spring, Volume 9, 4, pp. 79-114.

OH, H. L., and SCHEUREN, F. J. (1978). "Multivariate Raking Ratio Estimation in the 1973 Exact Match Study," Proceedings of the Section on Survey Research Methods, American Statistical Association, 107-111.

PROJECTOR, D. S. AND WEISS, G. S. (1966). Survey of Financial Characteristics of Consumers, Board of Governors of the Federal Reserve System, Washington, DC.

RAO, J. N. K., and WU, C. F. J. (1988). "Resampling Inference With Complex Survey Data," Journal of the American Statistical Association, 83, 401, pp. 231-241.

RUBIN, D. B. (1987). "Multiple Imputation for Nonresponse in Surveys," New York, John Wiley and Sons.

SKINNER, C. J., HOLT, D., and SMITH, T. M. F. (editors) (1989). Analysis of Complex Surveys, John Wiley and Sons.

WILSON, O., and SMITH, W. J. Jr. (1983). "Access to Tax Records for Statistical Purposes," Proceedings of the Section on Survey Research Methods, American Statistical Association, pp. 595-601.

WOODBURN, R. L. (1991). "Using Auxiliary Information to Investigate Nonresponse Bias," Proceedings of the Section on Survey Research Methods, American Statistical Association.