

**THE SURVEY OF INCOME AND
PROGRAM PARTICIPATION**

**LONGITUDINAL ANALYSIS OF
FEDERAL SURVEY DATA**

No. 112 9609

**P. Ruggles
Joint Economic Committee**

June 1990

ACKNOWLEDGEMENTS

Any opinions expressed in this paper are the author's alone and should not be attributed to the Joint Economic Committee or the Bureau of the Census.

TABLE OF CONTENTS

I.	Introduction.....	1
II.	Creating a Longitudinal Analysis File.....	3
	Choosing the Unit of Analysis.....	4
	Constructing a Person-Level Longitudinally	
	Linked File.....	8
	Dealing with Multi-Wave Data.....	12
	Attrition and Longitudinal Weighting.....	15
	Imputation in the Longitudinal Context.....	17
III.	Making Comparisons Across Time.....	19
	Comparing Two Points in Time.....	19
	Examining Transition Events.....	20
	Analyzing Data on Duration.....	23
	Other Issues in Longitudinal Analysis.....	29
IV.	Conclusions.....	31

REFERENCES

I. Introduction

Longitudinal panel data provide a unique opportunity to examine patterns and sources of economic and demographic change at the individual and family level. These data are relevant to a host of policy issues, from the assessment of welfare program participation to an understanding of patterns of health care usage or of the determinants of retirement. Many policy issues require some understanding of the factors that lead up to a particular event, or of the consequences that stem from it. Without repeated observations of the individuals concerned, however, such factors and consequences can only be inferred. Thus, our increasing store of longitudinal panel data holds the potential for major breakthroughs in our understanding of the basic determinants of economic and demographic change as they affect individuals and families over time.

Unfortunately, however, many of our longitudinal data sets have been somewhat under-used by researchers so far, especially compared to similar cross-sectional surveys. To some extent this under-usage may simply stem from the fact that many of these data sets are still fairly new--researchers need a chance to become familiar with the opportunities offered by these new sources of information. A more fundamental problem, however, is that to an analyst whose primary research experience is with cross-sectional microdata, a longitudinal panel of microdata on families and individuals can be rather intimidating.

A longitudinal database designed to offer a reasonably representative sample of the non-institutionalized population, for example, will be much larger and more complex than a similarly representative cross-sectional sample, since every observation will have been repeated several times. Size alone is likely to create some analysis problems, and the mechanisms used by the panel's designers to track

individuals and their various relationships over time are also likely to complicate the process of analysis. To produce readily interpretable results from such panels, the analyst must think carefully about the unit of analysis to be examined and about the construction of appropriate summary measures across demographic units and across time.

Further, problems such as reporting errors that would be hidden in a cross-sectional survey such as the Current Population Survey (CPS) are likely to be disconcertingly obvious in a longitudinal survey--we know, for example, that individuals very rarely change genders, and it is also quite unusual for one's age to decline. Such problems can shake the analyst's faith in the quality of the data set as a whole, even if only a few such errors appear. Additionally, any longitudinal survey is likely to experience some attrition over time, leading to difficult questions on the appropriate use of imputations and on the construction of sample weights for longitudinal analysis. While there is a fairly large literature on these issues as well as on the specific techniques that can be used in analyzing longitudinal data, the application of each of these issues to the particular problem to be analyzed must be considered at least briefly by any researcher undertaking a longitudinal analysis project.

The purpose of this paper is to provide some guidance to users and potential users of longitudinal data sets who are trying to sort out appropriate approaches to the problems of analyzing longitudinal panel data. This paper does not attempt to offer any new insights into the methodologies available to estimate the determinants of change (or stability) in a given variable or set of variables over time, nor are the theoretical issues underlying these methodologies addressed in any detail. Instead, the paper is designed to be a much more basic "how to" guide, focusing on the most fundamental choices that must be made by the analyst in undertaking a project

involving the use of longitudinal data. Further, the paper focuses almost exclusively on the application of longitudinal analysis to questions concerning patterns of family income, expenditures and/or demographic change.

The remainder of this paper is organized into three sections. The first of these addresses basic issues in designing a file for longitudinal analysis. The most crucial of these issues, in my view, is choosing the appropriate unit of analysis for the application at hand. This section discusses the pros and cons of alternative choices, and considers the implications of these choices in constructing an analysis file. Other problems in file construction--dealing with multi-wave data, handling attrition bias and longitudinal weighting, and the pros and cons of various types of imputation--are also considered very briefly.

The following section considers specific methods of making comparisons across time. The major focus of this section is on matching the outcome measures and statistical techniques chosen to the basic research question being asked. For many policy issues fairly simple outcome measures may be perfectly appropriate, but it is important to understand the measurement implications of alternative choices in order to avoid misinterpreting one's results.

The final section of the paper describes a few specific examples of current approaches to the measurement of economic and demographic phenomena using longitudinal microdata. It then concludes with a discussion of additional steps that federal statistical agencies could take to facilitate the analysis of the various longitudinal databases they produce.

II. Creating a Longitudinal Analysis File

Most longitudinal data on individuals, families and households come from surveys that consist of a series of interviews with selected sample members

conducted at more or less regular intervals. Examples of longitudinal surveys that are constructed in this way include the Panel Study of Income Dynamics (PSID), the Survey of Income and Program Participation (SIPP), the National Longitudinal Surveys of Labor Market Experience (NLS), the Retirement History Survey (RHS), and the household survey portions of the National Medical Care Expenditures Survey (NMCES) and the National Medical Care Expenditures and Utilization Survey (NMCUES).¹

In many cases, longitudinal data are released first in the form of cross-sectional files pertaining to each set of interviews, and individual analysts may be required to link the files across time themselves before they can be used longitudinally. Even when these linkages are provided in some form by the data collectors, however, analysts must typically decide how to carry forward information about household structure and family relationships from one interview to the next. In other words, the analyst must choose the basic unit of analysis around which the longitudinal file is to be organized.

Choosing the Unit of Analysis. Choosing the unit of analysis for longitudinal research poses some problems that generally do not arise in cross-sectional analysis. Many of our most familiar cross-sectional data sets--for example, the CPS--are arranged hierarchically, with information on several different possible units of analysis presented together. Thus, a typical CPS record will include information on the household, followed by information on any families or sub-families in the household, and those data in turn will be followed by the individual person-records

¹This list is by no means exhaustive, but these are some of the surveys most commonly used by policy analysts interested in economic and demographic issues. This paper focuses on longitudinal applications in these areas, and, more narrowly still, on applications that deal specifically with microdata on families and individuals. For a broader discussion of available longitudinal data sets, see Office of Management and Budget, Statistical Policy Office (1986).

for each household member. It is a fairly simple matter to analyze data across all households, across families or sub-families, or across individuals--or even to combine information from all three levels into one analysis.

Although even in the cross-sectional case the choice of a unit of analysis will have some impact on one's outcome measures--the measured poverty rate for families, for example, is different from the rate for individuals--definitions of the terms "household," "family," and "person" are by now familiar, and the implications of choosing one unit rather than another are generally clear. The appropriate unit of analysis in such a case will typically depend on the purposes of the analysis being undertaken, rather than on the constraints imposed by the data.

Unfortunately, even relatively straightforward terms such as "household" or "family" lose a great deal of their precision when they are considered longitudinally, however. Although one may think of the family income reported in the CPS as applying to essentially the same "family" over the period of a year, for example, in fact families may undergo a substantial amount of change over the course of a year. This problem is addressed in the CPS by fixing the composition of the family at the point of the interview. Retrospective data is then collected on the incomes over the previous year of all those who happen to be in the family on the interview date, regardless of whether each specific person was a member of that family for the entire year.

This approach can be duplicated using longitudinal survey data, but for a survey like the SIPP which offers actual month by month data on both family composition and income such an approach seems both cumbersome and potentially

misleading.² Indeed, since many of the policy and research issues that are of interest in a longitudinal context involve the relationship between changes in family composition and changes in income, using this methodology could undermine the advantages associated with longitudinal analysis as a whole.

Unless family or household composition is fixed arbitrarily at some point in time, however, longitudinal analysis at the family or household level requires a new set of definitions of what it means to be a family or household. In other words, one must decide how much change is acceptable within the limits of the term "family." If a couple has a baby is that a new family, or a continuation of the old one? If they get divorced, is one of the new units to be treated as a continuation of the old family, or are they both new? What if an elderly parent of one member moves in? Is the resulting family the continuation of both previous families, of one of them, or of neither? Should the answer to this question depend on who is designated to be the head of household in this new arrangement?

Considerable analysis of the implications of alternative family definitions in a longitudinal context has been carried out by Connie Citro and Roger Herriot (among others).³ They have addressed most of the questions raised above, and have worked out the empirical implications of specific alternatives that involve different decision rules for continuing or terminating someone's membership in a given family. Summarizing broadly, they find that the specifics of a longitudinal family definition matter somewhat in considering issues such as changes in poverty status,

²A survey like the PSID that collects data annually rather than monthly will typically use a fixed family composition within any given year. If these data are used longitudinally, however, the analyst's attempts to trace income and family composition over a period of years will result in unit-definition problems similar to those described here in relation to monthly data.

³See for example Citro et al. (1986) or McMillen and Herriot (1984) for further discussion and for additional references on this point.

and the use of some longitudinal definition produces results that are quite significantly different from those seen when family composition is treated as fixed.

The idea of using a longitudinal family definition is attractive, at least initially, to many policy analysts--after all, many, even most, policy issues of interest pertain to the family, not to the individual. Indeed, income considered at the person level is a fairly meaningless concept, since a major assumption of both our income support policies and of our social support system in general is that the members of families and households pool their incomes, and at least to some extent, make joint consumption decisions.

Over the past several years, however, both the work on this topic by Citro, Herriot, and others and my own experiences in analyzing longitudinal income data from the SIPP and the PSID have convinced me that the use of a longitudinal family concept can in fact be quite misleading. In general, in order for a measurement concept to be useful, minor variations in its specification should not result in major differences in the quantities being measured. When measures are not robust in this way, it is difficult to tell whether specific outcomes are related to actual differences in behavior or some other factor across groups, or whether they are simply artifacts of the measurement method. In my experience, longitudinal concepts of the family typically fail this test. Additionally, use of such a concept can actually impede longitudinal analysis, since the very factors that are of interest for much policy research--the impacts of divorce, out of wedlock births, deaths, and so forth--also tend to change family composition, and under many definitions result in new families. Linking these new families with their predecessors in a way that facilitates our understanding of the impacts of these transitions can become very difficult, since families can combine and recombine in many different ways over a given observation period.

As a result, therefore, longitudinal linkages at the person level are generally the most satisfactory solution to the problem of linking records across time. After examining a great deal of longitudinal data on family structures and incomes, it has become apparent that the person is the only unit of observation that is really reasonably constant over time. While people do move in and out of the sample over a period of time, the identification of a specific individual during his or her time in the sample can almost always be made unambiguously.

A second reason for preferring person-level linkages is that, again after some data-analysis experience, I have also become convinced that counts of the number of people with a given characteristic or behavior pattern are generally at least as useful for policy purposes as are counts of families or other aggregate units. If we are able to estimate that 50,000 people will enter poverty as a result of a given legislative proposal, for example, we may not much care if those people are in 15,000 families or 25,000. Indeed, if two proposals could be shown to affect the same number of people but different numbers of families (if one particularly targeted large families, for example), most analysts would consider it misleading to cite only the number of families involved.

Constructing a Person-Level Longitudinally Linked File. Nevertheless, person-level linkages do have some important disadvantages. In order to work with person-level data, it is necessary to find some way around the problems that originally led Citro, Herriot et al. to try and develop a longitudinal family concept. Specifically, for example, personal income (without additional family income information) does not provide a very good estimate of the resources actually available to each person. Additionally, most transitions of interest involve changes in the family as a whole, not just in specific members. For these reasons, substantial modifications must be made in the simple person-level linked longitudinal file if it is to be used for policy

analysis. For the most part, these modifications involve moving additional information on the family and household in which each person resided at each observation point onto the person record.⁴ In essence, this will create a personal longitudinal family history for each person in the longitudinal file.

To illustrate, consider the case of a married couple who, within the observation period contained in the longitudinal file, have a baby and then become divorced. Further, suppose that both adults were working at the beginning of the period, but that the woman left her job when the baby was born, and the man then experienced a spell of unemployment that lasted until after the divorce. Clearly, in order to have a complete picture of the events that have happened in this family it is necessary to have some information on the family as a whole--neither adult's person-record alone will contain all the necessary information (such as, for example, the other adult's work status).

On the other hand, a longitudinal family record will probably not contain all the information either--after the divorce, one or both of the adults will normally be considered part of a new family. As a result, that person (or both persons) will appear in two separate families within the observation period, making it difficult to examine issues such as the impact of the divorce on each spouse's income and poverty status. If either spouse remarries (or even moves in with a parent or other

⁴The slightly awkward term "observation point" is used here to refer to the unit of time over which specific longitudinal surveys collect their observations. This unit will typically be either a month (as in the SIPP, the NMCUES and the NMCES) or a year (as in the PSID and most cohorts of the NLS). A few surveys (e.g. the RHS) collect their data over longer or even irregular intervals. Note that the "observation point" is not necessarily the same as the "interview point," since some surveys collect information on several intervals of time that have occurred between interviews. As discussed briefly below, this technique is likely to lead to irregularities in the reported data series, but may facilitate using the data to simulate "event histories"--continuous records of the duration of specific states--which is helpful in applying certain analytic techniques.

relative) his or her family picture will change again, again creating discontinuities in the total record if a longitudinal family definition is used.

While in theory it might be possible to maintain linked records on every family that contained common members within the observation period, in practice this becomes extremely cumbersome in any file that contains data from more than three or four observation or interview dates. Among other problems, this approach requires a complex system of pointers to indicate which groups of people actually belong together at any given observation point, and where each of them are at the previous and subsequent observation points. For many longitudinal applications, therefore, this approach is impractical.

The alternative is to move a subset of the family information for each observation date on to the individual person records, and then to link these records at the person level. To consider cases like the one outlined above, for example, one might move information on total family income, spouse's earnings and employment status, family size, and number and ages of children at each observation date onto each adult's person-record for that date. This would allow one to observe changes in income and family composition as they affect each member of the household over the observation period as a whole--information on the previous family status of the non-custodial parent would not be separated from his or her record in the period after the divorce, for example.

The specific family variables suggested above for inclusion on the person-record are not the only ones that might be used, of course--for other types of analyses we have included markers for variables such as the presence (and/or age) of an elderly person in the family, receipt of specific types of income such as welfare benefits by someone in the family, whether the family is female-headed at a given observation date, and so forth. Many of these are fairly obvious, and will be used

by many analysts--examples include family income and family size. Such very common variables should almost certainly be constructed and appended to every adult's person-record for every observation date by the agency issuing the longitudinal data set. For family and household variables of importance to a specific analysis that are not routinely coded onto the person-record, however, the analyst must create a specific routine to produce appropriate recodes. In the longer run, it would again be helpful if agencies producing longitudinal data files could provide analysts with some help in implementing such recodes.

In addition to such simple variables, which can be extracted directly from the family record and appended to the person record, it is also very helpful for longitudinal analysis to create some specific transition flags that relate to the family changes under examination. For example, in considering the impacts of divorce on income and poverty status it would be helpful to have flags indicating changes in marital status. Appending such a flag to the person-record for the observation point in which the divorce takes place greatly simplifies the examination of the impacts of this change on other variables.

In fact, certain transitions are so important in determining income and family status that the data-issuing agencies should be encouraged to flag them routinely before issuing longitudinal data files. Such transitions might include divorce, marriage, birth of a child to a family member, death of a family member, and the loss or gain of a job by a family member.

Once family variables and transition flags have been appended to each person-record, it is a fairly straightforward matter to examine changes in family status and their impacts on income. To consider the impacts of divorce, for example, one would simply examine all the person records that contained the divorce flag. Family income and family size for each such person at each observation point could be

calculated directly from the person-record, with no additional linkages needed. Income could be summed over observation points, as necessary, to obtain an annual income measure. If desired, that measure could be compared to a similar sum of the per-period poverty standards for the various-sized families the person has been part of, to obtain a measure of the person's poverty status over a specified period before or after the divorce. Because the date of the divorce is flagged, it is relatively easy to compare pre- and post-divorce incomes over persons with different characteristics and with families of different sizes and types.

Because each person's family size may vary across observation points, it is helpful to use either the poverty standard or some other equivalence scale in considering relative incomes over a period of time--for example, before and after a divorce. Such an approach produces a slightly more difficult to interpret income variable (expressed as a percentage of the poverty standard, for example.) Nevertheless, this approach is considerably more enlightening in analytic terms than, say, assuming a fixed family composition over time, as cross-sectional surveys such as the CPS do.

Dealing with Multi-Wave Data. For budgetary reasons, longitudinal surveys often collect data at longer intervals than their designers would like. For example, surveys that ask about monthly income, family status, or medical care usage may collect information only every three or four months--coming back every month just isn't practical, and could also alienate respondents. Besides, many things don't change all that much from month to month, so that, given budget constraints, it is often preferable to put any extra dollars into an expanded sample size, for example, rather than more frequent interviews. Similarly, surveys designed to be administered annually occasionally get postponed, resulting in a two-year interval

between interviews, although those interviews may nonetheless ask about annual income and family data.

Surveys of this type, where the time-unit observed and the interview schedule are not coincident, pose some problems for the analyst. Inevitably, respondents are most likely to report any changes in income or family status as having occurred either at the beginning or the end of the interview period. Changes in income and other transitions, therefore, appear to occur much more heavily at the "seams"--the months (or other intervals) that represent the end of one interview period and the beginning of the next. This results in a situation where the probability of a transition is overstated at the beginning and end of the interview period, and understated in the intervening interval.

The extent to which this is a problem for the analyst depends on the type of analysis being done. If one is primarily interested in the determinants of a particular type of transition, for example, a slight misreporting of the date of the transition may not matter much--the respondent who misremembers when a given event took place is also likely to misremember other related events, and the relationship between events may therefore not be distorted at all. To put it another way, a respondent reporting both a divorce and a job loss is likely to get them in the right sequence--"I got divorced right after I lost my job"--even if he misreports the date of one or both events.⁵ Similarly, if the topic of interest is an ongoing state at the time of a particular transition--"I was unmarried and unemployed when I first went on welfare"--it is also likely to be reported correctly even if the date of the transition itself is slightly off.

⁵Nathan Young has examined the correlation of events on and off the seam, and has indeed found that in most cases correlations remain similar even when reported rates of transition are very different in the interview month and in other months. See Young (1989) for more discussion.

A pattern of small numbers of transitions occurring in the "off-seam" months followed by perhaps two to three times as many transitions in the "seam" month can be quite disconcerting, and can particularly make one doubt the usefulness of longitudinal data for duration analysis--questions like, "how long does a person typically stay unemployed after losing a job?" or "how long is the average welfare spell?" Indeed, such a pattern does distort reported spell lengths somewhat, creating more spells with durations equal to the interview interval, and fewer spells of other lengths. This distortion is less pronounced than is the distortion in transitions, however, since even for those who misreport one transition, it is not necessarily the case that both the onset and the termination of the spell will be reported as occurring on the seam.

Further, although the pattern of spell lengths resulting from data of this type undoubtedly contains some distortions, it is not clear whether the median spell length will typically be increased or decreased as the result of such misreporting. Indeed, overall one might expect the median spell duration to be largely unaffected, since there is generally no reason to believe that people under-report spell lengths more than they over-report them, or vice versa. In other words, for everyone who reports his current spell as having started a month earlier than it really did, there is probably someone else reporting a transition a month late. Unless the intervals between interview dates are very long, therefore, for many analytic purposes such multi-wave data, although disconcerting to look at, may still be a reasonable proxy for true event-history data.⁶

⁶In cases where the intervals between interviews are very long--anything much over a year is probably doubtful--treating the resulting data as if it were event history data could be quite misleading, since it is likely that many transitions could be missed altogether. This approach is particularly problematic if the survey in question was not designed to be continuous, and therefore does not ask specifically about changes occurring at shorter intervals within the interview period (the RHS,

Attrition and Longitudinal Weighting. A problem related to "seam-bias" is that most longitudinal surveys lose some respondents over the course of their operation, and unfortunately these losses are often correlated with the very factors that the analyst would like to study. Indeed, some apparent transitions at the seam are in fact caused by people dropping out of the survey, since a failure to participate in the next interview is of course always discovered at the scheduled interview date. This type of seam bias can be eliminated by carefully distinguishing between those who report the termination of a given state and those who leave the sample, as discussed further in section III below, but the underlying problem of attrition is more difficult to solve.

Problems of attrition are more important for some types of analysis than for others. For comparisons of repeated cross-sections, for example, they are very important--it is easy to mistake changes in the sample for changes in the underlying population if there is differential attrition across sample sub-categories. Similarly, any analysis that attempts to describe the incidence of a given type of transition over time may be vulnerable to this problem. In many cases, it will be appropriate for the analyst to standardize across the population eligible to experience the transition in question, although sometimes that population cannot be defined narrowly enough to eliminate the problem of differential attrition, which may be correlated with unobserved variables.

The problem of attrition may have less impact on duration analyses, unless the spells being examined tend to be long relative to the observation period, in which case it may be difficult to find a reasonably representative sample of completed

for example, is not designed for use as a continuous survey.) As discussed further below, for such surveys analysis as repeated cross-sections is appropriate, but duration-related analyses probably are not.

spells. Similarly, if those with long spells are particularly likely to leave the sample, the apparent duration of spells will tend to be biased downward. This is particularly likely to occur where the spells in question involve extended periods of low income, which in turn tends to be correlated with sample attrition.

Finally, it is unlikely that the problems described above can be completely solved through the use of appropriate longitudinal weights.⁷ A common approach to longitudinal weighting is to eliminate all part-panel cases, and then weight the remaining cases up to some control totals for the sample universe as a whole. Because a panel covers a certain period of time and people's circumstances change over time, these control totals must represent the population at some point in time--typically, the beginning of the panel. But, to the extent that those who leave the sample are indeed different from those who remain, simply weighting the remaining population so that they resemble the starting sample will not take care of the problem--those who remain will still behave differently, over the life of the panel, than those who have left.⁸

Indeed, for duration analysis at least, it is almost certainly preferable to retain part-panel participants for that period in which they are observed, in order to minimize the effects of attrition bias as much as possible. In this case, if one is not

⁷The following discussion of longitudinal weighting and imputation is at best a brief introduction to these subjects, and is directed to users and potential users of existing panel data, particularly the SIPP, the PSID, the NMCUES and the NMCES. For more details on the methods used to create longitudinal weights for existing surveys and on some possible alternative weighting strategies, see for example Cox and Cohen (1985), Ernst et al. (1984), Jones (1982), Kasprzyk and Kalton (1983), and Whitmore et al. (1982).

⁸If we could perfectly observe the characteristics of stayers and leavers, we could probably find some sample of stayers that otherwise behaved like leavers and could then reweight the sample appropriately. In fact, however, statistical agencies tend to rely heavily on basic demographic data such as age, sex, and race in designing longitudinal weighting cells, and these cells do a very imperfect job of capturing the differences between stayers and leavers.

using a self-weighted sample weights reflecting their relationship to the original sample universe should be used for all panel members, not just those in the sample for the full panel. For comparing repeated cross-sections, on the other hand, weights specific to the date of each cross-section should be used.

Imputation in the Longitudinal Context. Another approach sometimes used by data producers to handle the problem of attrition, as well as other reporting and response problems, is to impute variables and sometimes even whole records. Such imputations must be approached with caution. On the one hand, they can be very helpful in the context of repeated cross-sections, allowing one to carry out meaningful analyses without cumbersome changes in weights for each new cross-sectional observation. Even in this context, it is important to satisfy oneself that the imputations have been done in a way that provides some reasonable assurance that the outcome variables of major interest have not been seriously distorted, but for many types of analysis reasonably good imputations are quite possible.

Unfortunately, however, in dealing with longitudinal analyses it is often the analyst's first task to sort through the data and remove all the cross-sectional imputations put in by the data producer. Because many longitudinal panels are produced first as cross-sectional files they may contain imputations that are reasonable in the cross-sectional context but that are not designed with longitudinal applications in mind. For example, income imputations may be done for each wave without regard to the individual's income in any other wave, producing strange income patterns if the data are analyzed over time. In the longer run, it would be helpful if data-producing agencies would avoid cross-sectional imputations that distort the longitudinal record altogether, although this does require the

development of more sophisticated longitudinal imputation procedures than those now generally in use.⁹

After removing any cross-sectional imputations added by the data producer, the analyst must decide whether or not to make additional longitudinal imputations. For example, should the data be edited to eliminate miscodes, such as changes in gender, that are obvious in the longitudinal context even if not in cross-section? Where a given income amount shifts back and forth in type across waves--say, from child support to welfare and back again--is the analyst justified in "correcting" the record?

In general, I would argue that where the analyst has some strong reason to believe, based on the longitudinal record, that specific data is misrecorded, he or she should edit the record. Similarly, where there is missing data--say an age or even an income amount--that can be guessed within reasonable limits based on data in the surrounding waves, I would impute the missing value. While this may sound heretical--surely analysts should not "make up" data!--one must remember that in the context of longitudinal analysis such missing or incorrect values may be interpreted as spell beginnings or endings when in fact they are not, and failing to edit them may create events that in fact haven't happened. Where the probability is very strong that a given variable is misrecorded, therefore, some editing of the longitudinal record may be appropriate. Again, in the longer run, of course, it would be very helpful if statistical agencies would take a more aggressive approach to longitudinal editing and imputation--perhaps shifting resources away from cross-

⁹This set of complaints particularly applies to the SIPP, although other surveys where waves are produced first for cross-sectional analysis and are only later incorporated into some sort of longitudinal file may also be vulnerable to the same sort of problems.

sectional imputation, at least for databases that are designed with longitudinal applications primarily in mind.

So far this paper has discussed the creation of a longitudinal analysis file in considerable detail (although unfortunately there remain many conceptual issues in file construction that have only been touched on briefly or that have been neglected altogether, and the practical programming problems involved have not been considered at all). As has been hinted several times, however, many of the specific decisions that must be made in putting together a longitudinal research file depend on the specific issues to be examined, and in particular, on the types of outcome measures to be used. The next section, therefore, goes on to discuss alternative approaches to comparisons over time.

III. Making Comparisons Across Time

The major purpose of a longitudinal research file is of course to facilitate the analysis of change over time. There are three major types of time-related analysis that are commonly carried out with such files, and there are some specific methodological issues that pertain to each.

Comparing Two Points in Time. The simplest type of time-related analysis--the comparison of data from two discrete points in time--does not actually require a complete longitudinal data file at all. The major advantage of this type of analysis is that it is relatively simple to implement and can often yield a great deal of useful information, particularly for questions that focus on rates of turnover in a specific variable. This method is very commonly used with many different longitudinal data sets--several examples of such analyses can be found for PSID data in the Institute for Social Research's volume of PSID research results entitled *Years of Poverty, Years of Plenty*, for example. Other examples include Alan Fox's study using RHS

data which examined income changes at retirement, and the SIPP-based study produced by Jack McNeil and his colleagues at the Census Bureaus that considered how many of those poor in 1984 were still poor in 1985.¹⁰

The major drawback of this method of making comparisons across time is that the outcome variables are sometimes quite sensitive to the specific time periods chosen for analysis, and there is no way for the analyst to determine this if only two points in time are examined. Further, such comparisons are valid measures of change among those who already have a given characteristic, but cannot be used to determine the distribution of durations of a particular state among all those who enter it. For example, using this method we can tell what the total remarriage rate for all divorced women is over a given period of time, but we cannot determine the average amount of time that women spend between marriages, because we do not know when those who were already divorced at the time of the first observation got divorced, and we have no distribution of remarriage probabilities by duration of divorce to use in forecasting future remarriage rates for those who have not yet remarried. Indeed, we cannot even determine if the remarriage rate is sensitive to the amount of time that has elapsed since the divorce. In other words, to the extent that the determinants of changes in state are themselves time-related, they may be difficult to observe if one must rely on simple "before and after" comparisons.

Examining Transition Events. A second approach to making comparisons across time, therefore, is to examine transitions between two states directly. By focusing on the transition itself one can more closely examine its association with other factors that may not be observable in a simple before and after comparison.

¹⁰See Duncan (ed.) (1984) and McNeil et al. (1988).

This is helpful both in considering the effects of the transition on other variables and in estimating a causative model of the determinants of the transition itself.

To illustrate this point, let us reconsider the analysis of divorce discussed briefly above. If the analyst is interested not only in the determinants of the divorce transition, but also in its impacts, a simple comparison of two points in time may be doubly misleading. For example, family income may dip temporarily at the time of divorce as the family changes from one household to two. Eventually, however, as the two households make post-divorce adjustments in employment and living arrangements, income is likely to recover at least somewhat. Estimates of the impact of the divorce on income and poverty status for the various family members may be quite sensitive to both the unit definition used to compute income (as discussed in the last section) and to the specific timing of two income observations compared to the divorce itself.

In a case like this, examination of income or poverty status over a longer period leading up to and then following the transition will give a better picture of its actual impacts. For this type of examination it is necessary to have a longitudinally linked file with the transition flagged, but if such a file is available a descriptive analysis of this type is quite straightforward to perform.¹¹ Similarly, the transition flags themselves can be used as explanatory variables in a larger model of change over time as it affects some other variable. The recent paper by Suzanne Bianchi and Edie McArthur on the impacts of marital disruptions on children's economic status illustrates a transition analysis of this type.¹²

¹¹Applications illustrating the use of this technique to analyze income change can be found in Ruggles and Williams (1986) and Williams and Ruggles (1987).

¹²See Bianchi and McArthur (1989).

Considering the determinants of a given transition is also facilitated by the availability of a linked longitudinal file. For example, probit-type regression models can be used to examine the probability that a given transition will take place, subject to the various other characteristics of the cases in question. In analyzing divorce, for example, one might want to consider the impacts of the spouses' employment statuses in the period before the divorce on the probability that they will become divorced. In other cases, a broader set of dependent variables may be necessary--those leaving a given state may have more than one alternative option. The work by Alan Gustman and Thomas Steinmeier on retirement probabilities as observed in the RHS offers a good example of a fairly complex application of this type of transition analysis.¹³

With a linked longitudinal file, the conditional probability of a given event such as divorce or retirement can be calculated fairly easily for specific population subgroups, and/or conditioned on specific events, using readily available software packages such as SAS. Again, however, such an approach can be misleading if the determinants of the transition in question are themselves time-related--if for example, the previous duration of the marriage or even the length of the unemployment spell are important determinants of the probability of divorce.¹⁴

These duration-related issues, then, are potentially problematic with either a straightforward comparison of data from two points in time or with a more sophisticated analysis of specific transitions. Although it is sometimes possible to

¹³See for example Gustman and Steinmeier (1986).

¹⁴Additionally, if rates of divorce are changing rapidly over time, the use of pooled data on transitions from a long-term sample such as the PSID may give misleading estimates of transition probabilities. See for example Tuma and Hannan (1984) for more discussion of this point.

shoehorn duration-related information into one's transition analysis--one could create separate dummy variables for short and long unemployment spells in the above example, for instance--this is a rather ad hoc approach that is likely to leave many unanswered questions. In addition, in many cases one is interested not only in the transition event itself, or even in its impact on other events, but also in the expected duration of the new state that it creates. One wishes to know, for example, how long someone who enters poverty may be expected to remain poor, or how long someone who loses a job may be expected to remain unemployed. Questions of this type require some type of duration analysis.

Analyzing Data on Duration. There are many possible approaches to questions of duration, and alternative approaches can produce quite different and even seemingly contradictory statistics. The confusion generally results from differences in the population to which the duration estimate applies. The two major possibilities are cohort-based estimates, which typically apply to all those observed in a given state at a point in time, and spell estimates, which apply to all those observed to enter the state within a given span of time.

To illustrate these possibilities, consider the case of welfare program participation. A point-in-time or cohort-based estimate of welfare durations will ask a question like "How long have those who are currently receiving welfare been on the program?" This question has been phrased retrospectively, but it can also be put in a prospective form: "How long are those currently on the program likely to remain on in the future?" In either case, the base population being considered is all those on the program at a given point in time. Such estimates are therefore relatively easy to line up with cross-sectional estimates of the total population on welfare, which are of necessity also point-in-time estimates. Estimates of this type are very useful for a number of purposes--for example, estimating the future costs

of the current welfare caseload (although obviously to get total costs one would also have to account for new welfare entrants).

One useful way to think about estimates of this type is as an examination of the experiences of a particular cohort--a group that all happened to be in a given state at a given point in time. The NLS, for example, is designed with just such applications in mind. It is possible to use these data to examine the subsequent experiences of several distinct demographic cohorts selected at specific points in time--teenagers, men nearing retirement, women in their middle years. It is even possible, with the new youth cohort, to link up families across generations, and to relate young women's experiences to those of their mothers, as Peter Gottschalk has done recently for welfare recipients, for example. A similar type of application using PSID data is Frank Levy's path-breaking 1977 paper on the "underclass," which traced the subsequent experiences of a cohort of those in poverty in 1967.¹⁵

Cohort-type analyses are very useful for many policy questions, but it is important to be aware of their limitations in applying them to policy analyses. Specifically, because they apply only to those in the state at a given time, such analyses are sometimes difficult to generalize to the population as a whole, or even to the experience of all those who may pass through the state over a period of time.

What a point-in-time estimate cannot do, in other words, is answer questions like "How long will a typical person entering welfare stay on the program?" Such a question refers not to the population on the program at a point in time, but rather to the population entering the program. Although that may seem like a subtle distinction, in fact these two populations are likely to be very different if there is any significant variation at all in spell durations within the population as a whole. Those

¹⁵See Gottschalk (1989) and Levy (1977).

who are on welfare at a point in time are likely to have much longer spell durations, on average, than the typical entrant, because those with longer spell durations are more likely to be in the welfare population at any particular point in time.¹⁶

To see this point, consider a very simple example. Suppose the population of interest consists of 13 people, one of whom is in the state under consideration for one year, and twelve of whom are in that state for one month each. Further suppose that these twelve one-month spells are distributed so that one occurs in every month of the year. At any given point in time, therefore, the total population in the state being considered will consist of two people, one who is in a one-month spell, and one who is in a twelve month spell. A point-in-time analysis conducted any time after the first month will therefore conclude that 50 percent of the observable population reports a spell of more than one month. An analysis based on all entrants observed during the year, however, will find that only one-thirteenth of the population reports a spell of more than one month. Clearly, if the reasons for these differences in estimates are not well understood, they could lead to very different conclusions about the prevalence of long spells.

Many of the most useful and interesting questions that can be addressed using a longitudinal database are questions that relate to duration. In any type of duration analysis, however, it is necessary to be sensitive to the issue of censoring. Inevitably, there will be some spells that start before the beginning of the observation period or that end after the panel has come to an end. Further, there

¹⁶Mary Jo Bane and David Ellwood's classic paper on poverty spells makes this point very well, and provides a good example of spell analysis as applied to the PSID. (See Bane and Ellwood (1986)). For a similar example using SIPP data, see Ruggles and Williams (1989). Other useful applications include the work by Pamela Farley Short and her colleagues on spells of Medicaid participation and Rebecca Blank's imaginative use of longitudinal data from the Seattle and Denver Income Maintenance Experiments to examine spells of welfare program participation. See Short et al. (1988) and Blank (1986).

will be some cases that join the panel with a spell already in progress or leave the panel before one has ended. These spells cannot simply be ignored, since of course longer spells are more likely than short ones to be censored and ignoring this problem will therefore produce biased estimates.

An alternative approach that unfortunately is fairly often used by analysts who have not completely thought through the problem of spell censoring is to mix together all one's observations over a given span of time, whether they apply to completed spells or to those that are only partially observed. This produces results that are confusing and even potentially misleading, since it is easy to misclassify spells that are only partially observed as short spells, producing misleading estimates of average spell durations.

The measure of the "persistently poor" produced by Duncan et al. using PSID data is an example of this approach, and illustrates some of its problems.¹⁷ In this study, the base population was defined as all those in the population during the ten year observation period--not just those in poverty in a particular year, as in Levy's study. Duncan et al. then defined the "persistently poor" as those poor for at least eight out of the ten years. They went on to calculate the proportion of the total population that was "persistently poor" simply by dividing the number of people observed in poverty for at least eight years by the total population observed.

The problem with this approach is that some people who are poor for less than eight years during the observation period are nevertheless in the midst of spells of poverty that will total eight or more years--but unfortunately some of those years happen to fall outside the observation period. Thus the true number of individuals in the sample who were actually poor for at least eight out of ten years (at least

¹⁷See Duncan et al. (1984).

some of which fell in the sample period) cannot be estimated using these data. Estimates of the proportion of those observed who experience long poverty spells will be understated, because some spells that appear short are in fact longer, but they simply haven't been completely measured. At the same time, however, because these estimates mix together people who were poor in different years, they also cannot be used to predict, say, what proportion of those poor in a given year will still be poor eight years later.

Many analysts cope with the problem of estimating spell durations when some observations are censored by using some sort of survival analysis technique. Under this methodology, a survival function for a given type of spell is estimated based on the cumulative distribution of observed spell durations. In other words, in order to compute the probability that a spell of welfare participation, for example, will end in its sixth month, conditional on its having lasted for the first five months, one must include all cases known to have lasted at least five full months, whether or not their eventual disposition is known.¹⁸

To put it in more technical terms, the survival function for welfare participation may be estimated by defining $F^*(t, X_t)$ as the cumulative distribution of time on welfare, with X_t defined as a vector of independent variables affecting welfare participation (which may or may not vary with time themselves) and F^* representing the results of the series of participation decisions made to time t . At any time t , then, $F^*(t, X_t)$ may be seen as representing the probability that the

¹⁸This discussion is aimed at the analyst trying to decide whether this approach is appropriate for the particular application he or she has in mind. Anyone attempting to implement such an analysis should of course review some of the more technical literature on this topic. Tuma and Hannan (1984) provide a good basic overview of these methods. In addition, the treatment in Allison (1982) may be helpful to analysts who are completely unfamiliar with event history analysis techniques.

duration of welfare for someone with the given X vector of characteristics is less than t . The survival function for spells of welfare participation is then simply the proportion of observable spells still in progress at time t --that is, $S(t, X_t) = 1 - F^*(t, X_t)$.

By including all spells--even those whose endings will eventually be unobserved--for as long as information on their status is available, systematic biases related to spell duration will be minimized. At the same time, censored spells are essentially treated as if they had the same distribution of durations as spells with otherwise similar characteristics whose endings are observed. Under this methodology, censored spells do not pull down the estimated median spell duration, for example, as they do when the problem of censoring is not recognized. It is worth noting, however, that this approach assumes that censored spells are not systematically different from uncensored spells (except in ways fully captured in the X vector of explanatory variables), and that spells that occur at the beginning of the observation period are not systematically different from those starting nearer the end. To the extent that external events--for example, legislative changes or changes in the state of the economy--affect spell durations over time, analysis techniques that pool spell observations across the period as a whole may be misleading.

This approach does allow the contribution of a variety of factors--either fixed (e.g., sex and race) or time-varying (e.g., employment status)--to the conditional probability of exit (or of survival) to be estimated--these factors are simply included in the X vector of explanatory variables described above. This approach is very popular as a general method of analyzing spell durations and their determinants, and models of this type can be implemented in SAS as well as in other easily-obtained statistical packages (although typically the analyst is required to assume some specific underlying form for the distribution of exit probabilities). Only data sets

that provide a reasonably continuous record for a reasonably large sample of individuals entering the state being examined can be used with this approach, however, which limits its usefulness with smaller or less focused data sets or those in which data has been collected in an intermittent pattern.

Other Issues in Longitudinal Analysis. The most crucial decisions to make in undertaking a longitudinal analysis of family income and demographic data clearly involve the major choices concerning the basic outcome measures described above. A number of other issues also arise in longitudinal analysis, however, particularly in making income comparisons over time. Some of these issues--"seam" bias, recall and coding errors, the role of imputation--have already been discussed briefly in the section of the paper on file creation, but another set of issues--those relating to the accounting period for income measurement--can also have big impacts on one's outcome measure, especially in working with a file such as the SIPP that provides data on income over a sub-annual period.

In a cross-sectional file such as the CPS one does not have any particular choice over the accounting period that is used--income information is collected on an annual basis, and that is the only way it can be analyzed. In the SIPP, however, information is collected on monthly incomes. This allows the possibility that it can be examined over every period from one month to 32 months, the length of the panel. As work by Roberton Williams and our own more recent work has clearly demonstrated, alternative accounting period choices can have very different implications for income measures such as poverty rates.¹⁹ More than one fourth of the panel has at least one month with an income below the poverty level during calendar year 1984, for example, while only about 6 percent are poor in every month

¹⁹See Williams (1985) and Ruggles and Williams (1988).

of that year. Similarly, only about 20 to 25 percent of those who have an entry into poverty (measured using monthly income) during the panel are poor when income is measured on an annual basis. Because individuals' incomes fluctuate considerably even within a one-year period, it is important to consider the impacts of one's accounting period choices in any longitudinal income analysis.

These accounting period issues are important not only in analyzing a database like SIPP that collects sub-annual information, but also in comparing results across surveys using different accounting periods. For example, Bane and Ellwood's analysis of the duration of welfare spells performed using PSID data found a median spell duration of about two years, while a similar analysis of SIPP data found a median duration of just under a year.²⁰ Most of the difference is accounted for by the fact that the PSID uses an annual accounting period, while the SIPP uses a monthly one. Thus, someone who received welfare from October 1984 through March 1985 would be counted as having a two-year spell--covering both 1984 and 1985--in the PSID, while such a person would be counted as having a six-month spell in the SIPP.

There is no one correct approach to the accounting period issue--rather, the analyst should attempt to match the accounting period examined to the question being considered. For analyzing a program like Aid to Families with Dependent Children (AFDC), for example, where eligibility, benefits, and participation are all determined on a monthly basis, a monthly accounting period will generally be preferable where adequate data are available. For broader analyses of income distribution and even poverty, some analysts may prefer a longer accounting period.

For example, some argue that since people are able to average their incomes

²⁰See Bane and Ellwood (1983) and Ruggles (1988).

over time to some extent--save money in good months, say, to tide themselves over bad ones--longer accounting periods such as a year may give a better picture of people's real level of resources than do shorter ones. On the other hand, where detailed data are available it may be more appropriate to examine other resources such as asset holdings to judge total resources, rather than assuming that such resources are available on average over the longer period even when the family lacks income in the short run.²¹

IV. Conclusions

In summary, the many new sources of longitudinal data on incomes and family structures that have become available in the last decade offer exciting research opportunities to the policy analyst, but they bring with them their own unique measurement problems. Because these data sources are both more complex and less familiar than are cross-sectional databases covering such topics, analyzing them can present some challenges. For analysts willing to address these challenges, however, there are useful solutions, and these data can be used to provide important new insights into the processes underlying economic and demographic change.

Indeed, as discussed briefly in the various examples of measurement problems and their solutions given throughout the paper, important applications of longitudinal analysis to policy issues have already been carried out in many areas. A few examples include Bane and Ellwood's analysis of poverty spells and of AFDC participation using the PSID; the work by Bruce Vavricek and Ralph Smith of the Congressional Budget Office on spells of unemployment insurance reciprocity as observed in the SIPP; several Social Security Administration-sponsored studies on

²¹See Chapters 5 and 7 in Ruggles (1990) for more discussion of this point.

retirement behavior as observed in the RHS; and Peter Gottshalk's work on intergenerational transmission of dependency as observed in the NLS. Projects are now underway to address a whole host of additional issues, including patterns of health insurance coverage, multiple program participation for low-income beneficiaries, and earnings and employment patterns for the working poor.

The work that has been done so far and the work that is now underway represent major advances in our understanding of these issues, but there is much further analysis that could be done with our existing longitudinal survey data. To some extent, this expansion will simply take time--analysts need to become more familiar both with the surveys themselves and with appropriate techniques for analyzing and interpreting these data. Already, however, there is beginning to be a large literature on the applications of duration analysis, in particular, to economic and demographic data, and this literature can only be expected to grow over the next several years as additional data become available and additional issues are explored.

What can statistical agencies, and data producers in particular, do to help the analyst undertaking this type of study? In my view, these agencies could support longitudinal analysis efforts in two major ways.

First, as noted throughout this paper and particularly in the section on creating the longitudinal file, data producers do not always produce files that are highly amenable to longitudinal analysis, even when such analysis is the primary mission of a particular data-collection effort. Understandably, when a new survey such as the SIPP comes out a great deal of effort is devoted to the early cross-sectional files, since analysts are anxious to see how these new data line up with data from familiar cross-sectional surveys. In addition, the early waves of any survey will be ready for analysis long before the survey itself has been completed and edited longitudinally,

and data producers are understandably anxious to get these first products to the users as fast as possible.

Once a survey has been in regular production for some period of time, however, it would make sense to lessen the emphasis on cross-sectional files and to increase efforts to produce reasonable longitudinal data in a reasonably timely fashion. We already have excellent cross-sectional data on family incomes and labor force status, and unless the survey in question is clearly adding to our store of available cross-sectional data on a particular topic, cross-sectional applications should receive less attention. In particular, the level of effort devoted to activities such as cross-sectional imputation that have no application in the longitudinal context should be reduced. Instead, greater research efforts should be devoted to continuing problems like longitudinal editing and the development of reasonable longitudinal imputation procedures.

The second way in which statistical agencies could support longitudinal analysis would be to undertake more of it themselves. Data producers typically publish at least some cross-sectional information from the files they produce, and in some cases--the CPS publications in the Census P-60 series, for example, come to mind--these tables themselves provide important information on which policy-makers come to rely. It ought to be possible for the Bureau of the Census and other data producers to publish similar information, but of a longitudinal nature, using the longitudinal databases that they now produce.

The assumptions underlying survival analyses might be difficult to explain in such a context, but basic information on the experience of a given cohort, for example, is fairly easy to explain and to interpret. For instance, one could look at how many of those becoming unemployed in a given period were still unemployed

one, two, or more months later; how many of those on welfare or in poverty at a given point in time were still in that state x months (or years) later; and so forth.

Similarly, one could examine the transitions between states more directly, along with the characteristics of those experiencing the transitions. One could ask, for example, what proportion of those leaving unemployment in a given year find jobs, and what proportion leave the labor force? Does it differ for men and women, blacks and whites, old and young workers? For that matter, one could ask who becomes unemployed, and how does the incidence differ by demographic characteristics? Or, for example, what about those who enter welfare programs in a given year--what is the incidence of entry for those in different categories? What happens to those who leave welfare in that year? Do they get married? Do they get jobs? How many of those gaining jobs are still employed six months later, or a year later? Similar questions could be asked about the incidence and impacts of many other transitions, from divorce to retirement to the birth of a child.

The longitudinal analysis issues outlined above represent only a small proportion of those that could be undertaken--but the point here is that there is a great deal of fairly straightforward longitudinal analysis that would be very helpful to policy-makers, and that is not now being done in any systematic way.

Some very useful reports have been issued, of course--for example, the Census Bureau's P-70 series includes some longitudinal analysis from the SIPP, although so far such applications have been relatively limited in both quantity and scope. Again, many of these surveys, especially the SIPP and the NMCES, are still fairly new, so perhaps it is not surprising that their producers have not yet developed a complete, systematic schedule of reports examining basic longitudinal issues. Nevertheless, devoting more attention to their own longitudinal analyses would probably be the most important step data producers could take to support this type of research, and

could also increase substantially the useful information that we are able to obtain from these surveys.

REFERENCES

- Allison, P.D. "Discrete-Time Methods for the Analysis of Event Histories," in S. Leinhardt (ed.), *Sociological Methodology 1982*, San Francisco: Jossey-Bass, 1982.
- Bane, Mary Jo and David T. Ellwood. "The Dynamics of Dependence: the Routes to Self-Sufficiency." Report prepared for the U.S. Department of Health and Human Services. Cambridge, Mass.: Harvard University, 1983.
- Bane, Mary Jo and David T. Ellwood. "Slipping Into and Out of Poverty: The Dynamics of Spells." *Journal of Human Resources*, Winter 1986, 21(1), pp. 1-23.
- Bianchi, Suzanne, and Edith McArthur. "Family Disruption and Economic Hardship: The Short-Run Picture for Children." Paper presented at the annual meeting of the Population Association of America, May 1989.
- Blank, Rebecca. "How Important is Welfare Dependence?" Working Paper No. 2026. Cambridge, Mass.: National Bureau of Economic Research, Sept. 1986.
- Citro, Constance F., Donald J. Hernandez, and Roger A. Herriot. "Longitudinal Household Concepts in SIPP: Preliminary Results." SIPP Working Paper Series No. 8611. Washington D.C.: U. S. Bureau of the Census, 1986.
- Cox, B. and S. Cohen. *Methodological Issues for Health Care Surveys*. New York: Marcel Dekker, 1985.
- Duncan, Greg J. (ed.). *Years of Poverty, Years of Plenty*. Ann Arbor, Mich.: Institute for Social Research, 1984.
- Duncan, Greg J., Richard D. Coe, and Martha S. Hill. "The Dynamics of Poverty," in G. Duncan, ed., (op. cit.) 1984, pp. 33-70.
- Ernst, L., D. Hubble, and D. Judkins. "Longitudinal Family and Household Estimation in SIPP." Proceedings of the Survey Research Methods Section. Washington D.C.: American Statistical Association, 1984.
- Fox, Alan. "Work Status and Income Change, 1968-72: Retirement History Study Preview." *Social Security Bulletin*, 1976.
- Gottschalk, Peter. "The Intergenerational Transmission of Welfare Participation: Facts and Possible Causes." Paper presented at the annual meeting of the Association for Public Policy Analysis and Management, November 1989.
- Gustman, Alan L. and Thomas L. Steinmeier. "A Structural Retirement Model." *Econometrica*, May 1986, pp. 555-584.
- Jones, B. "Development of Sample Weights for the National Household Component of the NMCUES." Research Triangle Park, NC: Research Triangle Institute, 1982.

- Kasprzyk, Daniel and Graham Kalton. "Longitudinal Weighting in the ISDP," in Martin David, ed., *Technical, Conceptual, and Administrative Lessons of the ISDP*. New York: Social Science Research Council, 1983.
- Levy, Frank. "How Big is the American Underclass?" Working Paper 0090-1. Washington, D.C.: The Urban Institute, 1977.
- McMillen, David B. and Roger A. Herriot. "Toward a Longitudinal Definition of Households." SIPP Working Paper Series No. 8402. Washington DC: U.S. Bureau of the Census, 1984.
- McNeil, John, Enrique Lamas and Cynthia Harpine. "Moving Into and Out of Poverty: Data from the First SIPP Panel File." *Proceedings of the Social Statistics Section*. Washington DC: American Statistical Association, 1988.
- Office of Management and Budget, Statistical Policy Office. *Federal Longitudinal Surveys*. Statistical Policy Working Paper No. 13. Washington DC: OMB, May 1986.
- Ruggles, Patricia. *Drawing the Line: Alternative Poverty Measures and Their Implications for Public Policy*. Washington DC: Urban Institute Press, 1990.
- Ruggles, Patricia. "Welfare Dependency and Its Causes: Determinants of the Duration of Welfare Spells." Paper presented at the annual meeting of the American Economic Association, Dec. 1988.
- Ruggles, Patricia and Roberton Williams. "Longitudinal Measures of Poverty: Accounting for Income and Assets Over Time." *Review of Income and Wealth*, Sept. 1989, 35(3), pp. 225-244.
- Ruggles, Patricia and Roberton Williams. "Transitions In and Out of Poverty." Paper presented at the annual meeting of the American Economic Association, Dec. 1986.
- Short, Pamela Farley, Joel C. Cantor, and Alan Monheit. "Dynamics of Medicaid Enrollment." *Inquiry*, Winter 1984, 25(4), pp. 504-516.
- Tuma, Nancy B. and Michael T. Hannan. *Social Dynamics: Models and Methods*. New York: Academic Press, 1984.
- Vavrichek, Bruce and Ralph E. Smith. *Family Incomes of Unemployment Insurance Recipients and the Implications for Extending Benefits*. Washington DC: Congressional Budget Office, 1990.
- Whitmore, R., B. Cox, and R. Folsom. "Family Unit Weighting Methodology for the National Household Survey Component of the NMCUES." Research Triangle Park, N.C.: Research Triangle Institute, 1982.
- Williams, Roberton. "Poverty Rates and Program Participation in the SIPP and the CPS." Paper presented at the annual meeting of the American Statistical Association, August 1986.

Williams, Roberton and Patricia Ruggles. "Determinants of Changes in Income Status and Welfare Program Participation." Paper presented at the annual meeting of the American Statistical Association, August 1987.

Young, Nathan. "Wave Seam Effects in the SIPP." Paper presented at the annual meeting of the American Statistical Association, August 1989.